

# SOLUTION OF NONLINEAR STOKES EQUATIONS DISCRETIZED BY HIGH-ORDER FINITE ELEMENTS ON NONCONFORMING AND ANISOTROPIC MESHES, WITH APPLICATION TO ICE SHEET DYNAMICS\*

TOBIN ISAAC<sup>†</sup> GEORG STADLER<sup>†‡</sup> AND OMAR GHATTAS<sup>†§</sup>

**Abstract.** Motivated by the need for efficient and accurate simulation of the dynamics of the polar ice sheets, we design high-order finite element discretizations and scalable solvers for the solution of nonlinear incompressible Stokes equations. In particular, we focus on power-law, shear thinning rheologies commonly used in modeling ice dynamics and other geophysical flows. We use nonconforming hexahedral meshes and the conforming inf-sup stable finite element velocity-pressure pairings  $\mathbb{Q}_k \times \mathbb{Q}_{k-2}^{\text{disc}}$  or  $\mathbb{Q}_k \times \mathbb{P}_{k-1}^{\text{disc}}$ , where  $k \geq 2$  is the polynomial order of the velocity space. To solve the nonlinear equations, we propose a Newton-Krylov method with a block upper triangular preconditioner for the linearized Stokes systems. The diagonal blocks of this preconditioner are sparse approximations of the (1,1)-block and of its Schur complement. The (1,1)-block is approximated using linear finite elements based on the nodes of the high-order discretization, and the application of its inverse is approximated using algebraic multigrid with an incomplete factorization smoother. This preconditioner is designed to be efficient on anisotropic meshes, which are necessary to match the high aspect ratio domains typical for ice sheets. As part of this work, we develop and make available extensions to two libraries—a hybrid meshing scheme for the `p4est` parallel adaptive mesh refinement library, and a modified smoothed aggregation scheme for PETSc—to improve their support for solving PDEs in high aspect ratio domains. In a comprehensive numerical study, we find that our solver yields fast convergence that is independent of the element aspect ratio, the occurrence of nonconforming interfaces, and of the mesh refinement, and that depends only weakly on the polynomial finite element order. We simulate the ice flow in a realistic description of the Antarctic ice sheet derived from field data, and study the parallel scalability of our solver for problems with up to 383 million unknowns.

**Key words.** Viscous incompressible flow, nonlinear Stokes equations, shear-thinning, high-order finite elements, preconditioning, multigrid, Newton-Krylov method, ice sheet modeling, Antarctic ice sheet.

**1. Introduction.** We design high-order finite element discretizations and scalable solvers for incompressible nonlinear Stokes equations describing creeping flows of power-law rheology fluids. Applications include ice sheet dynamics [31], mantle convection [53], magma dynamics [44] and other problems involving non-Newtonian fluids [26]. Among the main challenges for the solution of these problems are the presence of local features that emerge from the nonlinear constitutive relation, the strongly varying and anisotropic coefficients arising upon linearization, the incompressibility condition leading to indefinite matrix problems, complex geometry and boundary conditions, a wide range of length scales that may require highly-adapted meshes with high aspect ratios, and large problem sizes that necessitate parallel solution on large supercomputers. Our approach to cope with these challenges uses adaptively

---

<sup>†</sup>Institute for Computational Engineering & Sciences, The University of Texas at Austin, Austin, TX, USA

<sup>‡</sup>Courant Institute of Mathematical Sciences, New York University, New York, NY, USA

<sup>§</sup>Department of Geological Sciences and Department of Mechanical Engineering, The University of Texas at Austin, Austin, TX, USA

\*Support for this work was provided by the U.S. Department of Energy Office of Science (DOE-SC), Advanced Scientific Computing Research (ASCR), Scientific Discovery through Advanced Computing (SciDAC) program, under award numbers DE-FG02-09ER25914, DE-11018096, and DE-FC02-13ER2612, and the U.S. National Science Foundation (NSF) Cyber-enabled Discovery and Innovation (CDI) program under awards CMS-1028889 and OPP-0941678. Allocations of computing time on TACC's Stampede under XSEDE, TG-DPP130002, and on ORNL's Titan, which is supported by the Office of Science of the DOE under DE-AC05-00OR22725 are gratefully acknowledged.

refined nonconforming meshes, high-order inf-sup stable finite elements, and iterative Newton-Krylov solvers combined with multilevel preconditioning techniques. We focus in particular on the construction of efficient solvers and preconditioners for discrete systems resulting from high-order discretizations.

High-order finite element methods for partial differential equations (PDEs) are attractive because, in many situations, the discrete solution converges rapidly to the true solution as the approximation order  $k$  is increased or the characteristic mesh size  $h$  is decreased. However, this increased accuracy per degree of freedom compared to low-order methods does not automatically translate into increased accuracy per unit of computational work. This is due to the fact that matrices arising from high-order discretizations are denser and, thus, more expensive to apply and to solve systems with. The cost of applying a matrix arising from a high-order discretization can be reduced drastically if the work is shifted from memory operations to floating point operations. This can be achieved using matrix-free implementations and tensor-product approximation spaces and element operations on hexahedral finite element meshes. To precondition matrices arising from high-order discretizations, low-order preconditioners based on the nodes of the high-order discretization have proven efficient [6, 11, 15, 29, 48]. These preconditioners allow fast construction and the use of methods established for low-order discretizations.

Our approach to solving the nonlinear Stokes equations is an inexact Newton-Krylov method, with a block preconditioning strategy for the linearized equations, built from preconditioners for the (1,1)-block and for its Schur complement. We consider a power-law rheology that involves the second invariant of the strain rate tensor, for which the Newton linearization results in a fourth-order anisotropic tensor viscosity. We pay particular attention to the interplay between discretization and solver: seemingly minor differences in either the discretization or the low-order preconditioner can vastly impact the performance of conventional solution methods for both diagonal blocks of the preconditioner.

Our driving application is the simulation of the dynamics of continental-scale ice flows, which is a critical component of coupled climate modeling. Predicting the contribution of ice sheets to sea-level rise is difficult because of the complexity of accurately modeling ice sheet dynamics for the full polar ice sheets and the large uncertainties in unobservable parameters governing these dynamics [45, Chapter 10, Appendix 6]. To address these uncertainties, significant effort has been focused on the development of inverse methods to infer ice sheet model parameters from observations [46, 50]. These inverse methods require the repeated solution of ice flow equations for numerous parameter fields that may vary over wide ranges, and many also require the repeated solution of related adjoint ice flow equations. Hence, inverse methods particularly stress the efficiency and robustness of solvers for nonlinear Stokes equations.

A particular difficulty in ice sheet simulations is the high aspect ratio of the computational domains, which is inherited by the discretization, leading to anisotropic meshes. Discretizations with high-aspect ratio elements (and problems with highly anisotropic material properties, which have many of the same properties) are known to be challenging for implicit solvers and preconditioners. The development of robust solvers for high aspect ratio domains is also important in other Earth science and climate modeling problems. In ocean flow models, for instance, three-dimensional implicit PDE models are now being used [36], whereas in the past they were often replaced by two-dimensional approximations.

*Related work.* Several recent articles develop scalable solvers for Stokes problems with varying viscosity [8, 10, 20, 22, 25, 27, 43]. These contributions use low-order stable or stabilized finite elements for the discretization of the Stokes equations and address nonlinearity mostly via a Picard fixed point approach. Scalable solvers for high-order discretizations of nonlinear scalar problems and extensions to linear incompressible flow problems are studied in [6]. Various asymptotics-based approximations of the Stokes equations are used for ice sheet and glacier modeling, which reduce the indefinite Stokes equations to positive definite elliptic systems. These simplifications are justified by the large differences between horizontal and vertical components of the velocity; we refer to [30] for a comparison and discussion of the validity of these different models. Ice sheet simulations using the full nonlinear Stokes equations can be found, e.g., in [23, 41, 49].

*Contributions.* One of the main contributions of this paper is the design of discretizations, solvers and preconditioners that allow the fast and scalable iterative solution of nonlinear Stokes problems. In particular, we obtain convergence that, for a large class of realistic problems, is independent of the mesh size, the presence of nonconforming interfaces in the mesh, and the element aspect ratio, and depends only weakly on the polynomial order. Another contribution is the extension of low-order preconditioners for high-order discretized problems to meshes with nonconforming interfaces and high aspect ratio elements.

In addition to analyzing our solver techniques on workstation-sized model problems, we also demonstrate their performance and scalability on a series of larger problems requiring a distributed memory parallel implementation, including a simulation of the dynamics of the Antarctic ice sheet. The simulation uses a geometry and temperature field derived from field data and constitutes what we believe to be the first highly resolved nonlinear Stokes-based continental scale simulations of the Antarctic ice sheet.

We have also developed publicly available tools for discretizing and solving PDEs on high aspect ratio domains, such as (but not limited to) those occurring for ice sheets. One, presented in section 3.1, is an extension to the `p4est` parallel adaptive mesh refinement (AMR) library [9] that allows it to construct favorable meshes for these domains. The other one, discussed in section 5.1, is a plugin for PETSc’s [4] generic algebraic multigrid (GAMG) preconditioner to improve its effectiveness for these types of discretizations.

*Limitations.* We use algebraic multigrid (AMG) for preconditioning. An alternative would be to use geometric multigrid (GMG), which builds a hierarchy through geometric coarsening of the mesh. GMG can be tailored to only coarsen in certain directions while leaving others unchanged (semicoarsening), which can be useful for anisotropic geometries [56]. We use the easier-to-use AMG, in which we incorporate a minimal amount of geometric information.

Our simulations of ice sheet dynamics use a fixed temperature field and geometry. Simulations of evolving ice sheets would require coupling of the nonlinear Stokes equations with a time-dependent advection-diffusion equation for the evolving temperature, and with a kinematic equation for the evolution of the ice sheet surface. However, the solvers presented in this paper carry over as important components in a time-stepping procedure for the simulation of time-evolving nonlinear viscous flows.

*Overview.* This paper is organized as follows. In section 2, we discuss the form of the nonlinear Stokes equations and boundary conditions that are the focus of this work, their variational formulation, and their linearization. In section 3, we present

stable mixed-space finite element discretizations and a discussion of adaptive mesh refinement. In section 4, we give an overview of our approach to solving the resulting discrete system of nonlinear equations and the linearized counterparts. The preconditioner for the (1,1)-block of the linearized Stokes systems that arise at each Newton iteration, which is a critical component of the linear solver, is presented in detail in section 5, followed by a discussion on the preconditioner for the Schur complement of the (1,1)-block in section 6. We test our solver on a model geometry in section 7, and then on a discretization of the Antarctic ice sheet in section 8, where we also study the scalability of our method. We conclude with a discussion in section 9.

**2. Nonlinear incompressible Stokes equations.** After specifying the Stokes equations with strain rate thinning power-law rheology in section 2.1, we present the corresponding variational form and argue existence of a unique solution in section 2.2. The linearization of the nonlinear equations is presented in section 2.3.

**2.1. Problem statement.** On an open, bounded domain  $\Omega \subset \mathbb{R}^3$  we consider the incompressible Stokes equations

$$-\nabla \cdot \boldsymbol{\sigma} = \mathbf{b}, \quad \mathbf{x} \in \Omega, \quad (2.1a)$$

$$\nabla \cdot \mathbf{u} = 0, \quad \mathbf{x} \in \Omega, \quad (2.1b)$$

where  $\mathbf{u}$  is the flow velocity and  $\mathbf{b}$  is a body force. The Cauchy stress tensor  $\boldsymbol{\sigma}$  depends on the strain rate tensor  $\mathbf{D}(\mathbf{u}) = \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^\top)$ , its second invariant  $\mathbf{D}_\Pi(\mathbf{u}) := \frac{1}{2}\mathbf{D}(\mathbf{u}) : \mathbf{D}(\mathbf{u})$  and, possibly, other physical quantities such as a temperature field. Here, “:” denotes the Frobenius product between second-order tensors  $\mathbf{A} = (\mathbf{A}_{ij})$  and  $\mathbf{B} = (\mathbf{B}_{ij})$  defined by  $\mathbf{A} : \mathbf{B} = \sum_{i,j} \mathbf{A}_{ij} \mathbf{B}_{ij}$ . In the ice sheet problem, which is our driving application, the stress tensor is given by Glen’s flow law

$$\boldsymbol{\sigma} = -p\mathbf{I} + B(T)(\mathbf{D}_\Pi(\mathbf{u}) + \varepsilon)^{\frac{1-n}{2n}} \mathbf{D}\mathbf{u}, \quad (2.2)$$

where  $p$  is the pressure,  $B(T)$  a positive-valued function of temperature  $T$ ,  $n \geq 1$  is the strain rate exponent, and  $\varepsilon > 0$  a small regularization parameter that prevents infinite effective viscosity for  $n > 1$ . For  $n = 1$ , eq. (2.2) reduces to a linear rheology, and it describes a strain-rate weakening non-Newtonian fluid for  $n > 1$ . A common value used for modeling the flow of glacial ice is  $n = 3$ . To complete the definition of the boundary value problem eq. (2.1), it remains to specify the boundary conditions.

In ice sheet simulations, different parts of the boundary require different combinations of Dirichlet, Neumann, and Robin-type boundary conditions. This makes ice flow a good problem for developing methods for other creeping flow problems with complicated boundary conditions. At the ice-air interface, the homogeneous Neumann condition  $\boldsymbol{\sigma} \mathbf{n} = \mathbf{0}$  holds. At the ice-water interface, the normal stress matches the hydrostatic water pressure, i.e.,  $\boldsymbol{\sigma} \mathbf{n} \cdot \mathbf{n} = -p_w$  and the tangential components of the boundary traction vanish. At the base of the ice sheet, complex interactions occur between ice, water, rock, and till. In cold regions, the ice sticks to the bedrock, while in temperate regions, water accumulates at the base and the ice can slide but is subject to some amount of friction. A general way to describe these phenomena is to use a Dirichlet condition in normal direction to describe melting and freezing at the base of the ice sheet, combined with a Robin-type sliding law relating the tangential component of velocity  $\mathbf{T}_\parallel \mathbf{u} = (\mathbf{I} - \mathbf{n} \otimes \mathbf{n})\mathbf{u}$  to the tangential component of the stress through a function  $\beta(\cdot, \cdot, \dots)$ , i.e.,

$$\mathbf{T}_\parallel \boldsymbol{\sigma} \mathbf{n} = -\beta(|\mathbf{T}_\parallel \mathbf{u}|, T, \dots). \quad (2.3)$$

Physically realistic descriptions of sliding must include the dependence of  $\beta$  on  $|\mathbf{T}_{\parallel} \mathbf{u}|$  [61], which makes eq. (2.3) a nonlinear boundary condition for the flow; in this work, however, we consider the linear case  $\beta(\mathbf{T}_{\parallel} \mathbf{u}) = \beta(\mathbf{x})\mathbf{T}_{\parallel} \mathbf{u}$ , and no basal freezing or melting. To summarize, we use the following boundary conditions for the base  $\Gamma_R$ :

$$\mathbf{T}_{\parallel} \boldsymbol{\sigma} \mathbf{n} + \beta(\mathbf{x})\mathbf{T}_{\parallel} \mathbf{u} = 0 \quad \mathbf{x} \in \Gamma_R, \quad (2.4a)$$

$$\mathbf{u} \cdot \mathbf{n} = 0 \quad \mathbf{x} \in \Gamma_R. \quad (2.4b)$$

Polar ice sheets have a characteristic depth of less than 5 kilometers, while they extend horizontally for thousands of kilometers. Because of this difference between length scales, modelers often simplify eq. (2.1) using asymptotic expansions that require assumptions about the magnitude of the velocities and stresses in the ice sheet, for instance, the shallow ice approximation [31] and the hydrostatic approximation [5]. The assumptions justifying these simplifications do not hold for the entire ice sheet, which has led to approaches that combine simplified models in the interior with Stokes equations at outlet glaciers [54]. To avoid these complications, we do not use simplified models and focus on the efficient solution of the Stokes equations (2.1) instead.

**2.2. Variational formulation.** Here, we define a variational form of eqs. (2.1) and (2.4) that defines the fields  $(\mathbf{u}, p)$  as the unique solution in a vector space  $\mathcal{V} \times \mathcal{M}$  to

$$\int_{\Omega} [\mu(\mathbf{u})\mathbf{D}(\mathbf{v}) : \mathbf{D}(\mathbf{u}) - p\nabla \cdot \mathbf{v} - q\nabla \cdot \mathbf{u}] \, d\mathbf{x} + \int_{\Gamma_R} \beta \mathbf{T}_{\parallel} \mathbf{v} \cdot \mathbf{T}_{\parallel} \mathbf{u} \, ds = \mathbf{f}(\mathbf{v}), \quad (2.5)$$

for all  $(\mathbf{v}, q) \in \mathcal{V} \times \mathcal{M}$ , where  $\mu(\mathbf{u}) = \mu(\mathbf{u}, T) = B(T)(\mathbf{D}_{\Pi}(\mathbf{u}) + \varepsilon)^{\frac{1-n}{2n}}$ , and  $\mathbf{f}$  is the sum of the effects of body and boundary forces. We assume that  $B(T) \in \mathcal{L}^{\infty}(\Omega)$  is uniformly bounded from below, that  $\beta \in \mathcal{L}^{\infty}(\Gamma_R)$  is nonnegative and that  $\partial\Omega$  is Lipschitz continuous. For simplicity, we assume that the Dirichlet boundary conditions are homogeneous and are incorporated into the space  $\mathcal{V}$ .

For a similar problem, Jouvét and Rappaz [35] show that a unique solution  $(\mathbf{u}, p)$  exists in the Dirichlet-conforming subspace of  $[\mathcal{W}^{1,r}(\Omega)]^3 \times \mathcal{L}^{r'}(\Omega)$ , where  $r = 1 + 1/n$  and  $r' = 1 + n$ . In appendix A we define a pair of spaces  $\mathcal{V}$  and  $\mathcal{M}$ , in which eq. (2.5) is well-posed. This pair is only slightly modified from the pair above to account for the linear Robin boundary condition.

**2.3. Newton linearization.** The Newton linearization of eq. (2.5) about a velocity-pressure pair  $(\mathbf{u}, p)$  are equations whose solution  $(\tilde{\mathbf{u}}, \tilde{p}) \in \mathcal{V} \times \mathcal{M}$  satisfies

$$\int_{\Omega} [\mathbf{D}(\mathbf{v}) : (\boldsymbol{\mu}'(\mathbf{u})\mathbf{D}(\tilde{\mathbf{u}})) - \tilde{p}\nabla \cdot \mathbf{v} - q\nabla \cdot \tilde{\mathbf{u}}] \, d\mathbf{x} + \int_{\Gamma_R} \beta \mathbf{T}_{\parallel} \mathbf{v} \cdot \mathbf{T}_{\parallel} \tilde{\mathbf{u}} \, ds = -\mathbf{r}(\mathbf{u}, p, \mathbf{v}, q) \quad (2.6)$$

for all  $(\mathbf{v}, q) \in \mathcal{V} \times \mathcal{M}$ . Here,  $\boldsymbol{\mu}'(\mathbf{u})$  is an anisotropic 4th-order tensor given by

$$\boldsymbol{\mu}'(\mathbf{u}) = \mu(\mathbf{u})\mathbf{I} + \frac{\partial\mu(\mathbf{u})}{\partial\mathbf{D}_{\Pi}(\mathbf{u})}\mathbf{D}(\mathbf{u}) \otimes \mathbf{D}(\mathbf{u}) = \mu(\mathbf{u}) \left( \mathbf{I} - \frac{n-1}{2n} \frac{\mathbf{D}(\mathbf{u}) \otimes \mathbf{D}(\mathbf{u})}{\mathbf{D}_{\Pi}(\mathbf{u}) + \varepsilon} \right), \quad (2.7)$$

and  $\mathbf{r}(\cdot)$  is the residual of eq. (2.5). Here, “ $\otimes$ ” denotes the outer product between two second-order tensors. Compared to the Newton linearization eq. (2.6), the commonly used Picard linearization of eq. (2.5) neglects the anisotropic part of the fourth-order tensor  $\boldsymbol{\mu}'(\mathbf{u})$ . Using a finite element discretization of eqs. (2.6) and (2.7) is only marginally more complex than the Picard linearization, as the action of  $\boldsymbol{\mu}'(\mathbf{u})$  on

$\mathbf{D}(\tilde{\mathbf{u}})$ , which is all that is required in a finite element method, can be computed using Frobenius products with trial and test functions. The operator  $\boldsymbol{\mu}'(\mathbf{u})$  is also found in the adjoint equations corresponding to eq. (2.5), which are used in inverse methods to infer uncertain parameters from observations [46, 50].

**3. Discretization.** Our goal for discretizing eqs. (2.5) and (2.6) is to obtain discrete solutions that converge to the continuous solutions rapidly in terms of the number of unknowns and time-to-solution. Our approach combines locally refined meshes with high-order finite element approximation spaces. The adaptive meshes we use are discussed in section 3.1, the finite element approximation spaces for the Stokes equations are described in section 3.2, and computational aspects of the discretization are addressed in section 3.3.

**3.1. Meshing.** We use a hierarchical approach to mesh refinement, which starts with a coarse mesh of conforming hexahedra. This coarse mesh is expected to roughly describe the geometry of the domain  $\Omega$ . The fine mesh used for the finite element discretization is obtained by hierarchical refinement of this coarse mesh. Refinement can be used to improve the resolution of the geometry of  $\Omega$  or the resulting small hexahedra can simply inherit the geometry. For practical as well as numerical reasons, we require our refined meshes to obey a 2:1 condition, where neighboring hexahedra can differ by only one level of refinement, as illustrated in fig. 3.2. We use the `p4est` library to manage refinement, to enforce this 2:1 condition and to partition the mesh between processes in parallel computations [9, 32]. We use an extension of the `p4est` library, developed by the first author, for meshing three-dimensional problems in nearly two-dimensional domains. This extension uses a forest-of-quadtrees to manage independently-refinable columns of hexahedra. Meshes created with this extension have three key properties that forest-of-octrees meshes lack.

- (1) Elements—and by extension degrees of freedom—are organized into columns that span the height of the mesh. This organization can be exploited by solvers: see section 5.1.
- (2) Each column of elements is assigned to a single process. Split columns, which can appear in meshes created by the forest-of-octrees approach, can have negative consequences for solvers: see section 5.1.
- (3) Hexahedra within a column may be locally refined in the vertical direction, allowing for more flexible refinement than the purely isotropic refinement of the forest-of-octrees approach. We exploit this fine control over the element aspect ratio in our meshes of the Antarctic ice sheet: see section 8.

This type of mesh refinement is illustrated in fig. 3.1. This extension to the `p4est` library appears as collection of “`p6est`” data types and functions, so named because they build on elements of both the “`p4est`” interface for forests-of-quadtrees and the “`p8est`” interface for forests-of-octrees. Documentation of the `p6est` interface is available online.<sup>1</sup>

Meshes for the simulation of ice flow must address the different length scales inherent in the problem. To accurately capture the vertical variations of the state variables, a minimum vertical mesh resolution is necessary. Most ice sheet models use  $\sim 10$  nodes in each vertical column and have a horizontal resolution of 5 km. Since the average thickness of the polar ice sheets is  $\sim 2$  km, the width-to-height aspect

<sup>1</sup>For documentation of `p6est`, see [http://p4est.github.io/api/group\\_p6est.html](http://p4est.github.io/api/group_p6est.html), as well as the `example/p6est/test/test_all.c` example distributed with the library, which demonstrates the major I/O, mesh refinement, and visualization functions.






FIG. 3.1: Illustration of the type of mesh refinement provided by the `p6est` extension of the `p4est` parallel AMR library. The hexahedral elements are organized into columns, which are always assigned to a single process when the mesh is partitioned. Note that horizontal refinement is decoupled from vertical refinement.

ratio  $\phi$  of these discretizations is  $\sim 25$ .<sup>2</sup> The Antarctic ice shelves, moreover, are typically  $\sim 500$  m thick, so the same horizontal resolution results in  $\phi \approx 100$ . We seek discretizations and methods for the Stokes equations that support these aspect ratios.

**3.2. Finite element discretization.** In this section, we describe the finite element spaces used to discretize the velocity and pressure spaces  $\mathcal{V}$  and  $\mathcal{M}$ .

**3.2.1. Discrete velocity spaces.** Given a mesh  $\mathcal{T}$  of hexahedra  $\{K_i\}$  with possibly nonconforming interfaces, we define a finite-dimensional subspace of  $\mathcal{C}^0(\Omega)$  using isoparametric Lagrange finite elements. The  $n_k := (k+1)^3$  nodes  $\Xi_k = \{\xi_{rst}\}_{0 \leq r,s,t \leq k}$  that define our Lagrange finite elements are the tensor-product Gauss-Lobatto nodes of polynomial degree  $k$  on the reference domain  $\hat{K} = [-1, 1]^3$ . These basis functions span  $\mathbb{Q}_k(\hat{K})$ , the space of functions on  $\hat{K}$  that are the univariate polynomials of degree at most  $k$  in each of the coordinate directions. We map  $\hat{K}$  to an element  $K_i$  by  $\varphi_i \in [\mathbb{Q}_k(\hat{K})]^3$  and use the tensor-product Gauss nodes of order  $k$  for numerical quadrature. We define the finite-dimensional space

$$V_{\mathcal{T},k} = \{v \in \mathcal{C}^0(\Omega) : \forall K_i \in \mathcal{T}, v \circ \varphi_i \in \mathbb{Q}_k(\hat{K})\},$$

and the velocity space  $W_{\mathcal{T},k} = [V_{\mathcal{T},k}]^3 \cap \mathcal{V}$ . For a conforming mesh, the set of element nodes  $\Xi_k$  naturally defines a set of global nodes  $X_{\mathcal{T},k}$  for  $V_{\mathcal{T},k}$  by the images of element nodes. For a nonconforming mesh  $\mathcal{T}$ , however, not all element nodes correspond to global nodes, as shown in fig. 3.2. To construct a set of global nodes at a nonconforming interface, we thus ignore some element nodes: if an element is more refined than its neighbor, the element nodes on the interface with that neighbor do not contribute to the set of points that define the global nodes. Instead, these nodes are known as hanging or dependent nodes.

In general, function values at the nodes of element  $K_i$  must be interpolated from the global vector of nodal values by a restriction matrix  $R_i$ . If  $K_i$  has no hanging nodes, then  $R_i$  is simply a one-to-one association of  $\Xi_k$  to a subset of  $X_{\mathcal{T},k}$ ; if  $K_i$  has hanging nodes, then  $R_i$  interpolates values as described in fig. 3.2. We use identical trial and test spaces, so a global nodal matrix  $A$  is assembled from element nodal matrices  $\{A_i\}$  by  $A = \sum_i R_i^\top A_i R_i$ .

**3.2.2. Discrete pressure spaces.** We use inf-sup stable mixed finite element spaces to avoid the artificial compressibility that can be introduced by stabilized discretizations of incompressible flow. Additionally, to satisfy element-wise incompress-

<sup>2</sup>In glaciology, one often uses the thickness-to-width aspect ratio  $\epsilon = \phi^{-1}$  as the relevant limit in asymptotic expansions is  $\epsilon \rightarrow 0$ . In this work, however, we prefer using  $\phi$  because we consider the thickness of an ice sheet to be its characteristic length, in which we also measure its horizontal extent.




FIG. 3.2: (a) A two-dimensional  $\mathbb{Q}_3$  mesh with a 2:1 nonconforming interface. (b) The nodal values along a nonconforming interface. Shown in green are the Gauss-Lobatto nodes of the smaller element, which do not align with those of the larger element: function values at these nodes depend on the values at the nodes of the larger element, so they are not included in the global nodal basis. The matrix  $R_i$  that interpolates a function to the nodes of the smaller element must interpolate cubic polynomials to the hanging nodes. This polynomial interpolation is dense: the value at each of the hanging nodes is dependent on all of the independent nodes. For two-dimensional nonconforming interfaces,  $R_i$  is defined by tensor-product polynomial interpolation.

ibility, we favor piecewise discontinuous pressure spaces  $\mathcal{M}_{\mathcal{T},k}$ . This mass conservation is particularly important for ice sheet simulations, where the change of the mass of the ice sheet is an important quantity of interest in climate projections.

The two most common choices for approximation on the reference cube are  $\mathbb{P}_{k-1}(\hat{K})$ , which is the space of polynomials on  $\hat{K}$  of degree at most  $(k-1)$ , and  $\mathbb{Q}_{k-2}(\hat{K})$ . We study two possibilities for inf-sup stable velocity-pressure finite element spaces. The pairing  $\mathbb{Q}_k(\hat{K}) \times \mathbb{P}_{k-1}(\hat{K})$  has an optimal order of convergence, and has an inf-sup constant that is independent of  $k$  and of the type of hierarchical local mesh refinement we use [28]. Its inf-sup stability, however, degrades with increasing aspect-ratio  $\phi$  [3]. We find (see section 6) that this degradation can be significant for the element aspect ratios in our meshes. An alternative pairing is  $\mathbb{Q}_k(\hat{K}) \times \mathbb{Q}_{k-2}(\hat{K})$ , which has a suboptimal order of convergence, but its inf-sup stability is uniform with respect to boundary layer refinement making it appropriate for large values of  $\phi$  [57]. For this pairing, the inf-sup constant decreases as  $\mathcal{O}(k^{-1})$ ; however, for the moderate values of  $k$  used in this work, this dependence is not problematic.

**3.3. Computational aspects.** For a mesh with  $N_{\text{el}}$  elements, the number of degrees of freedom  $N_{\text{dof}}$  in a  $k$ -order finite element discretization is  $\mathcal{O}(N_{\text{el}}k^3)$ , and the number of nonzero entries in the matrix for a system of equations defined on that space is  $\mathcal{O}(N_{\text{el}}k^6) = \mathcal{O}(N_{\text{dof}}k^3)$ . This increasing density means that sparse matrix-vector products based on globally assembled element matrices are not efficient (in terms of memory operations) for large values of  $k$ . We therefore compute nonlinear residuals and apply linear operators using a matrix-free approach to finite elements, where only the coefficients and fields that define an operation are stored in memory and the operation's application to a specific vector is assembled from all element contributions at the time of application. This approach requires  $\mathcal{O}(N_{\text{dof}})$  memory operations per matrix-vector product or residual calculation. For high-order elements, it is thus better suited to modern computer architectures, where the bandwidth for memory operations is much narrower than the bandwidth for floating-point operations. This reduction in memory operations comes at the expense of more floating point operations, but the tensor structure of  $\mathbb{Q}_k(\hat{K})$  allows for all such element computations to be reduced to repeated applications of one-dimensional compute kernels, which can



be highly optimized [6, 16]. The same one-dimensional kernels are used to apply the restriction operators  $R_i$  for hanging nodes.

**4. Newton-Krylov method for nonlinear Stokes equations.** Our goal is to design a robust and scalable solver for the nonlinear Stokes equations eqs. (2.1) and (2.2) with boundary conditions that include eq. (2.4). Ideally, the convergence should be stable with respect to: (1) the element size and the mesh refinement pattern, (2) parameters in the rheology  $\mu(\mathbf{u})$ , (3) the Robin coefficient field  $\beta$ , (4) the polynomial order  $k$ , and (5) the element aspect ratio  $\phi$ . Here, we propose an inexact Newton-Krylov solver, outlined in sections 4.1 and 4.2 below. In section 4.3, we present a test problem that allows us to study the effects of the factors listed above. Variations of this problem will be used in sections 5 and 6 to analyze and optimize the convergence of our linear solver, and in section 7 for a nonlinear Stokes problem.

**4.1. Newton’s method for nonlinear Stokes equations.** Given a velocity and pressure pair  $(\mathbf{u}, p)$ , we (approximately) solve eq. (2.6) for a search direction  $(\tilde{\mathbf{u}}, \tilde{p})$ . We then conduct a line search in the direction  $(\tilde{\mathbf{u}}, \tilde{p})$  using the weak Wolf conditions [47] to guarantee that the nonlinear residual decreases. Each Newton update is computed inexactly via a Krylov-space iterative method, but with a tolerance that decreases in subsequent steps so as to guarantee quadratic convergence of the Newton iterations [18] close to the solution. We ensure that discretization and differentiation commute, so that the Jacobian obtained from discretizing eq. (2.6) is equivalent to differentiating the discretization of eq. (2.5). Discretization with one of the stable finite element pairs discussed in section 3.2 results in a linear system with the symmetric saddle-point system matrix

$$A(\mathbf{u}) = \begin{pmatrix} F(\mathbf{u}) & B^\top \\ B & 0 \end{pmatrix}, \quad (4.1)$$

where the (1,1)-block  $F(\mathbf{u})$  is the discretization of the sum of the terms involving  $\beta$  and  $\mu'(\mathbf{u})$  in eqs. (2.6) and (2.7), and  $B$  is the discretized divergence operator.

**4.2. Preconditioned Krylov method for linearized Stokes equations.** We solve systems involving  $A(\mathbf{u})$  using preconditioned Krylov space methods, typically restarted GMRES, or, if the preconditioner is not constant, its flexible variant FGMRES [51]. As is well known, the performance of Krylov methods critically depends on the availability of an efficient preconditioner  $\tilde{A}$  for  $A(\mathbf{u})$ . In the following, we use the notation  $A = A(\mathbf{u})$  and  $F = F(\mathbf{u})$ , i.e., in our notation we neglect the dependence of  $F$  and  $A$  on  $\mathbf{u}$ . Due to the elliptic nature of  $F$ , a purely local preconditioner for  $A$  cannot provide  $h$ -independent convergence and a multilevel preconditioner is required. There are two main approaches for multilevel preconditioners for incompressible flow problems, namely monolithic and block preconditioning approaches. The former approximates the saddle point system on each level of a hierarchy and employs smoothers that are based on approximate local saddle point solutions (i.e., Vanka-type smoothers) [14, 34]. This approach typically requires a geometric mesh hierarchy or involves stabilized discretizations. In contrast, block preconditioners are built from preconditioners for  $F$  and for the Schur complement with respect to the (1,1)-block,  $S := -BF^{-1}B^\top$ . They allow to build on existing solvers for elliptic systems and do not impose restrictions on the discretization underlying  $A$ . Due to this flexibility, we follow this latter approach and use an upper-triangular block preconditioner

tioner  $\tilde{A}$ , such that the preconditioned system becomes

$$A\tilde{A}^{-1} = \begin{pmatrix} F & B^\top \\ B & 0 \end{pmatrix} \begin{pmatrix} \tilde{F} & B^\top \\ 0 & \tilde{S} \end{pmatrix}^{-1} = \begin{pmatrix} I + \varepsilon_{\tilde{F}} & \varepsilon_{\tilde{F}} B^\top \tilde{S}^{-1} \\ B\tilde{F}^{-1} & I + \varepsilon_{\tilde{S}} \end{pmatrix}, \quad (4.2)$$

where  $\varepsilon_{\tilde{F}} = I - F\tilde{F}^{-1}$  and  $\varepsilon_{\tilde{S}} = I - (-B\tilde{F}^{-1}B^\top)\tilde{S}^{-1}$ . Here, the matrix  $\tilde{F}$  is an approximation of the (1,1)-block  $F$ , and  $\tilde{S}$  is an approximation to the Schur complement  $S$  with respect to the (1,1)-block. If  $\varepsilon_{\tilde{F}}$  and  $\varepsilon_{\tilde{S}}$  vanished,  $A\tilde{A}^{-1}$  would be a lower triangular matrix with all eigenvalues clustered at 1, and a preconditioned Krylov method would converge in two iterations. Thus, our target is to devise  $\tilde{F}^{-1}$  and  $\tilde{S}^{-1}$  such that  $\varepsilon_{\tilde{F}}$  and  $\varepsilon_{\tilde{S}}$  are small and  $\tilde{A}$ -preconditioned Krylov methods converge quickly. Additionally, it is important that the setup time for the preconditioner  $\tilde{A}$  is small, because  $\tilde{A}$  is recomputed as the Jacobian  $F = F(\mathbf{u})$  changes. In sections 5 and 6, we develop block preconditioners  $\tilde{F}$  and  $\tilde{S}$  and study their properties.

In sections 7 and 8, when we apply our preconditioner to problems with dimensional coefficients, we combine the preconditioner  $\tilde{A}$  that we describe in this section with a diagonal rescaling of the linearized Stokes system of the type described in [43, Section 2.6]. This rescaling tends to reduce the round-off errors that result from matrix entries that differ by several orders of magnitude due to different units.

**4.3. Test problem setup.** Our test problems are based on problem C from the Ice Sheet Model Intercomparison Project [49]:  $\Omega$  is a cutout of an “infinite slab”, i.e., a sheet that is periodic in  $x$ - and  $y$ -directions. We use homogeneous Neumann boundary conditions on the top surface of  $\Omega$ ; on its base, we employ homogeneous Dirichlet conditions in the normal direction and Robin-type conditions in the tangential directions, as in eq. (2.4). We use the same boundary conditions when testing systems in just the (1,1)-block  $F$ . In linear test problems, we use  $\mu = 1$  and  $\beta = 1$  for the material and the sliding coefficients, unless specified otherwise.

To study the behavior of our solvers in the presence of nonconforming faces, we use meshes for  $\Omega$  that have nonconforming faces throughout, as illustrated in fig. 4.1. We use two base meshes:  $\mathcal{T}_{xy}$ , which has nonconforming interfaces in only the  $x$ - and  $y$ -directions, and  $\mathcal{T}_{xyz}$ , which has nonconforming interfaces in all directions. Note that the latter mesh cannot be decomposed into columns, so it represents a mode of mesh refinement that appears when using octree-based meshes, but not when using the hybrid AMR approach discussed in section 3.1. In all test problems,  $\Omega$  has unit thickness, but we vary the length and width  $L$  of the domain. The finite elements stretch with the domain, so that the number and pattern of elements is the same, but the element aspect ratio  $\phi$  varies between 1 and 100; see fig. 4.1b. These meshes have 576 elements, and the number of degrees of freedom in the Stokes systems discretized on these meshes is roughly  $2300(k-1)k^2$ .

For linear test problems, we test the effectiveness of our methods when the residual contains multiple length scales. To achieve this, we compute right hand sides from manufactured solutions, i.e.,  $b = A(\mathbf{u}^*)(\mathbf{u}^*, p^*)^\top$  when testing the linear Stokes solver and  $b = F(\mathbf{u}^*)\mathbf{u}^*$  when testing the (1,1)-block solver. For that purpose, we create scalar-valued, spatially variable fields  $s$  as the sum of a Fourier series with random coefficients and pointwise random component:

$$s(x, y, z) = \sum_{\substack{j,k=0 \\ (j,k) \neq (0,0)}}^N (a^{j,k}, b^{j,k}, c^{j,k}, d^{j,k})^\top \begin{pmatrix} \cos \omega j x \cos \omega k y \\ \cos \omega j x \sin \omega k y \\ \sin \omega j x \cos \omega k y \\ \sin \omega j x \sin \omega k y \end{pmatrix} |(j, k)|^{-\gamma} + e(x, y, z), \quad (4.3)$$




FIG. 4.1: Meshes and manufactured solution used for the test problems. Elements are colored according to their level of refinement. (a) Mesh  $\mathcal{T}_{xy}$ , which contains nonconforming faces in  $x$ - and  $y$ -directions, for  $\phi = 1$ . (b) Mesh  $\mathcal{T}_{xy}$  for  $\phi = 10$ . (c) Mesh  $\mathcal{T}_{xyz}$ , which contains nonconforming faces in all directions. (d) A manufactured solution  $u^*$ , constructed to contain variations with several length scales.

where  $N = 10$ ,  $\omega = 2\pi/L$ ,  $\gamma = 3/2$ , and  $i = 1, 2, 3$ . The coefficients  $a^{j,k}$ ,  $b^{j,k}$ ,  $c^{j,k}$  and  $d^{j,k}$  are randomly chosen from  $[-1, 1]$ , but the Fourier coefficients decay because of the  $|(j, k)|^{-\gamma} := (j^2 + k^2)^{-\gamma/2}$  term, making  $\gamma$  a control of the smoothness of  $s$ . The extra term  $e$  is a random value from  $[-1/4, 1/4]$ , added at each node of the discrete vector. To generate the vector field  $\mathbf{u}^*$  that is the manufactured solution of our test problems for the (1,1)-block solver, we generate a field  $s$  for each component of  $\mathbf{u}^*$ . The magnitude of such a velocity field is shown in fig. 4.1d. In all our tests, the norm we use to report the convergence of iterative solvers is the  $\ell_2$ -norm of the residual.

**5. Preconditioning the (1,1)-block.** The (1,1)-block  $F$  occurring in the Stokes system is similar to the operator arising in linear elasticity. If we neglect boundary conditions, its nullspace is given by the rigid-body modes. To approximately apply the inverse of  $F$ , we use algebraic multigrid (AMG) and, in particular, we use the smoothed aggregation multigrid (SA-AMG), which has theoretically proven convergence bounds [59]. SA-AMG uses Galerkin projections to create coarse approximation spaces, where the coarse space is embedded in the fine space by a prolongation matrix  $P$ .  $P$  is constructed by first creating a projector  $\tilde{P}$  that projects coarse “aggregate” nodes onto disjoint sets of fine nodes, followed by creating  $P$  from  $\tilde{P}$  by applying a local smoothing operation based on  $F$ , while ensuring that the near nullspace of  $F$  (the nullspace in the absence of boundary conditions) is well-approximated in the coarse space [60]. To construct the prolongation and coarse matrices of the hierarchy, we use SA-AMG as implemented by PETSc’s GAMG preconditioner. To build smoothers for each level, we use PETSc’s KSP and PC objects for defining Krylov methods and preconditioners. In this framework, we have three main design parameters, namely:

- (1) The matrix  $\tilde{F}$  used to construct the multigrid hierarchy. As multigrid is only used as preconditioner and not as solver,  $\tilde{F}$  can be based on a lower-order

(and thus sparser) approximations of the high-order element discretization of the (1,1)-block used in the residual computation.

- (2) How the hierarchy is coarsened. If fine nodes are only grouped into aggregates when they are strongly connected to each other, then there will be more aggregates and a less aggressively coarsened hierarchy. In standard SA-AMG, connections between fine nodes are ignored when forming aggregates if they are weaker than some threshold  $\theta$ .
- (3) The choice of smoothing operation on each level of the hierarchy.

These three choices cannot be made independently. The aggregation strategy and the smoothers are related because the error components not damped by the smoother must be corrected on the coarser levels. Moreover, the choice of  $\bar{F}$  affects the hierarchy and the effectiveness of different smoothers. In this section, we study these parameters on linear systems involving the (1,1)-block  $F$  rather than the full Stokes system  $A$ .

We note that two additional parameters, the size of the coarsest grid and the solver used for the corresponding system, can also impact the scalability and performance of multigrid preconditioners. In this work, however, we have used the default coarse grid size of GAMG (which is less than 50 unknowns), and we use a direct solver for the corresponding system in all of our numerical experiments. Note that in parallel simulations, we repartition the coarsened multigrid levels to subsets of processors, and, in particular, the system for the coarsest grid only uses a single processor.

**5.1. SA-AMG aggregation and smoothers.** The convergence of multigrid using pointwise smoothers such as Jacobi and Gauss-Seidel is known to degrade for anisotropic problems when isotropic coarsening is used [58, Chapter 4]. This slowdown can also be seen in fig. 5.1a, where, for different element aspect ratios  $\phi$ , we show the convergence based on a symmetric Gauss-Seidel (e.g., symmetric successive over-relaxation, SSOR) smoother for hierarchies created with aggregation threshold  $\theta = 0$ . This behavior occurs because for large  $\phi$ , horizontal variations in the error have less energy in the operator norm than vertical variations, and are thus not reduced sufficiently. There are two primary approaches to addressing this issue. One can use pointwise smoothers and use smaller aggregates to construct an operator hierarchy such that the undamped error components are well represented on the coarser meshes and can be corrected there (semicoarsening), or one can continue to use full-sized aggregates and use a smoother that more effectively dampens all error components (non-pointwise smoothing).

Although some form of semicoarsening could be effective for our problems, we have not had success with it. In standard SA-AMG, semicoarsening is accomplished by only aggregating strongly connected nodes, for which the corresponding matrix entries satisfy

$$|F_{ij}|^2 \geq \theta |F_{ii}| |F_{jj}|. \quad (5.1)$$

While this heuristic is generally good for scalar elliptic problems, for the vector system  $F$  it is problematic. It is known that for a fixed threshold  $\theta > 0$ , a sufficiently large aspect ratio  $\phi$  will cause the standard SA-AMG algorithm to break down for operators whose nullspace contains the rigid-body modes [24, Section 5.3]. Even before this breakdown, we observed poor convergence rates with  $\theta > 0$  and pointwise smoothers. A better heuristic for anisotropic problems than eq. (5.1), which is reported to result in good convergence for anisotropic elasticity problems, can be found in [37, Section 4], but even in this case, semicoarsening produces larger coarse operators and deeper hierarchies, which increase the cost of storing and applying the preconditioner. We

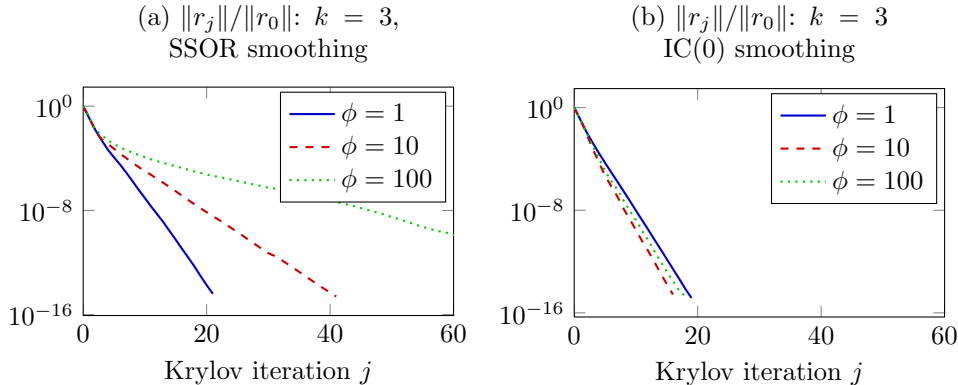


FIG. 5.1: Convergence histories for the  $(1,1)$ -block test problem described in section 4.3, discretized with  $k = 3$  on mesh  $\mathcal{T}_{xy}$  (fig. 4.1a) at different aspect ratios. The solver is GMRES preconditioned by one SA-AMG V-cycle. The solver in (a) uses a Chebyshev(1)/block-Jacobi/SSOR smoother, i.e., one optimally-damped application of block-Jacobi/SSOR; the solver in (b) is the same, except that it uses an incomplete Cholesky factorization (IC(0)) instead of SSOR.

therefore choose not to use thresholding when constructing our hierarchies, i.e., we do not consider the magnitudes of matrix entries when constructing our aggregates, only the nonzero pattern of the matrix. This choice results in aggressively coarsened hierarchies with smaller operator complexities, and has the additional advantage that it allows us to reuse the projection matrices for multiple Newton steps, reducing the cost of the subsequent preconditioner setups.

Since we must select a smoother that is compatible with our aggressive coarsening, we use approaches based on incomplete factorizations of the matrix  $\tilde{F}$ . In parallel, we only compute the incomplete factorization for each process’s diagonal block of the distributed  $\tilde{F}$  matrix, which amounts to a block-Jacobi smoother. To build a stationary smoother from this block-Jacobi/IC(0) smoother, an estimate of the largest eigenvalue of the smoothed operator is required, either to calculate a single damping parameter or the coefficients of a Chebyshev polynomial. We numerically estimate the largest eigenvalue using an iterative method, which adds to the setup cost of the multigrid preconditioner. To avoid this setup cost, one can alternatively use a non-stationary smoother that does not require damping, such as a few iterations of a Krylov method. This is a good choice when only a small reduction in the residual is required, as for instance in the early iterations of an inexact Newton-Krylov method.

We note that in certain parameter regimes—large values of  $\beta$  or no-slip Dirichlet conditions, and a fixed horizontal resolution—incomplete block-Jacobi factorizations have been observed to work well as stand-alone preconditioners, even with fine vertical resolution [7, Section 3.1]. However, we are not aware of a systematic study of the parameter regimes where this good performance can be observed.

Using local Fourier analysis, one can show that geometric multigrid smoothed by incomplete factorization can give  $\phi$ -independent convergence for scalar elliptic problems [58, Chapter 7]. A key component of this analysis is the ordering of degrees of freedom in the incomplete factorization: tightly-coupled degrees of freedom within a column must be sequential. We can order the degrees of freedom in our original system this way, but the standard techniques for generating aggregates, such as GAMG’s

(a) standard MIS					(b) column-preserving MIS				
$\ell$	$\beta = 1^0$	$1^{-2}$	$1^{-4}$	$1^{-8}$	$\ell$	$\beta = 1^0$	$1^{-2}$	$1^{-4}$	$1^{-8}$
0	0.14	0.14	0.57	0.63	0	0.14	0.14	0.30	0.35
1	0.17	0.27	0.75	0.78	1	0.17	0.20	0.39	0.47
2	0.20	0.51	0.82	0.83	2	0.19	0.25	0.48	0.54

FIG. 5.2: The convergence factor for SA-AMG applied to the (1,1)-block test problem for increasing mesh refinement and decreasing boundary condition strength. The aspect ratio of the elements is  $\phi = 100$ ; the polynomial elements order is  $k = 3$ ; the meshes are defined by  $\ell$  levels of isotropic refinement of the mesh  $\mathcal{T}_{xy}$ . We show the convergence factor (estimated over the number of iterations needed to reduce the residual by a factor of  $10^{-14}$  or 100 iterations) for decreasing magnitudes of the basal Robin coefficient  $\beta$ . Table (a) shows the results for a hierarchy constructed using standard randomized MIS aggregate construction, and table (b) shows the same for a hierarchy constructed using the column-preserving MIS aggregate construction described in the text.

default method of randomized maximal independent set (MIS) selection [2], result in coarse discretizations having no special spatial structure.

Because of this, the multigrid hierarchies created by randomized MIS aggregation are efficient for solving the (1,1)-block equations at a high aspect ratio only if the boundary conditions are strong enough, which we demonstrate in fig. 5.2a. If the basal boundary has full Dirichlet conditions, or a strong Robin coefficient  $\beta$  ( $\beta \geq 1$  in our non-dimensional model problem), we see that GMRES preconditioned by such a V-cycle converges with near  $h$ -independence. If, however, we have  $\beta \approx 0$ , then the V-cycles become less efficient and lose  $h$ -independence. The culprits for this loss of efficiency are error modes which are highly oscillatory in the  $x$ - and  $y$ -directions but nearly constant in the  $z$ -direction. At high aspect ratios, the stress energy of these modes is low in the absence of a strong boundary condition, so the smoother does little to dampen them. If the coarse mesh correction is poor, then it will introduce error components in these modes.

Our solution to this problem of weak boundary conditions is to use meshes whose degrees of freedom form columns such as the hybrid AMR meshes discussed in section 3.1, and to modify the aggregation technique used by GAMG in order to preserve the column-structure of these degrees of freedom. We first construct aggregates of columns using a standard aggregation technique such as randomized MIS, and then partition each aggregated column into the final node aggregates. When subdividing a column, we try to ensure that each node aggregate is at least three nodes tall, because three is usually the minimum diameter of aggregates used in standard SA-AMG. The multigrid hierarchies created by this method are much closer to achieving  $h$ -independent convergence for weak boundary conditions, as we demonstrate in fig. 5.2b. We note that this method appears to be effective even when the number of degrees of freedom in each column varies as they do in our test meshes. We also note that this technique does not depend on our particular choice of hexahedral elements, but could also be applied, e.g., to meshes with triangular prism elements, which are commonly used in modeling ice sheets and other high aspect ratio domains. We implement this aggregation technique in the DofColumns plugin<sup>3</sup> for GAMG.

<sup>3</sup>The DofColumns plugin is available from <https://bitbucket.org/tisaac/dofcolumns/>.



**5.2. Construction of low-order approximation  $\tilde{F}$ .** AMG requires assembled matrices for the construction of the coarse grid hierarchy. For matrices arising from high-order discretizations, this assembly requires significant memory and computation compared to low-order discretizations [6, 29]. In three dimensions, the cost to construct an element matrix  $F_i$  is  $\mathcal{O}(k^7)$  per element, or in terms of the number of nodes per element  $\mathcal{O}(n_k^{7/3})$ . Instead of the true element matrix  $F_i$  for element  $K_i$ , we construct an approximation  $\tilde{F}_i$  that treats the variables associated with nodes of the  $\mathbb{Q}_k(\hat{K})$  finite element as variables for a  $k \times k \times k$  grid of  $\mathbb{Q}_1(\hat{K})$  finite elements with corners at the high-order node locations. This matrix  $\tilde{F}_i$  can be constructed in  $\mathcal{O}(n_k)$  steps.

Aspects of this lower-order preconditioning technique have been studied for simple problems: spectral equivalence between high-order and lower-order discretizations of the Laplacian has been proven in two dimensions [39] and is demonstrated numerically in three dimensions. The (1,1)-block  $F$  differs from the operators used in previous studies in that it involves variable coefficients, high-aspect ratio elements, and non-conforming interfaces between elements. As we will show below, the interaction of these factors affects the stability and effectiveness of low-order preconditioning.

**5.2.1. Influence of quadrature on low-order preconditioning.** The effectiveness of the AMG preconditioner for  $F$  also depends on the choice of the quadrature used to construct  $\tilde{F}$ . Let us denote by  $\tilde{F}_G$  and  $\tilde{F}_{GL}$  the low-order matrices computed with Gauss-Legendre and Gauss-Lobatto quadrature, respectively. We find that  $\tilde{F}_{GL}$  leads to a better and more robust preconditioner, particularly when combined with an incomplete factorization smoother. The greater stability of Gauss-Lobatto quadrature comes from its diagonal-lumping behavior [17]. In tensor-product elements, the directional derivative operators that are used to construct the element matrices can be decomposed as Kronecker products of the form  $B \times B \times D$ , where  $B$  maps 1D nodal values to 1D values at quadrature points and  $D$  maps 1D nodal values to derivatives at quadrature points. When we refer to Gauss-Lobatto quadrature as “diagonal-lumping”, we mean that for Gauss-Legendre quadrature,  $B$  has off-diagonal entries, whereas for Gauss-Lobatto quadrature  $B$  is the identity. This does not make  $\tilde{F}_{GL}$  diagonal, but does increase sparsity:  $\tilde{F}_{GL}$  contains  $\sim 30\%$  fewer nonzero entries than  $\tilde{F}_G$ . The factors of an incomplete factorization can become ill-conditioned when the original matrix is far from being diagonally dominant, particularly if there is no pivoting performed during the factorization [52, Chapter 10.5]. The diagonal-lumping inherent in Gauss-Lobatto quadrature increases the magnitude of diagonal entries relative to off-diagonal entries. An additional advantage of the sparsity of  $\tilde{F}_{GL}$  relative to  $\tilde{F}_G$  is the reduced cost of the incomplete factorization and the hierarchy construction.

**5.2.2. Influence of nonconforming meshes on low-order preconditioning.** As discussed in section 3.2,  $C^0$ -conforming discretizations on nonconforming meshes require element restriction matrices  $R_i$  that contain dense blocks for nonconforming interfaces. For an element with hanging nodes, the product  $R_i^\top \tilde{F}_i R_i$  can at best be computed in  $\mathcal{O}(k^5)$  time, or in terms of  $n_k$ ,  $\mathcal{O}(n_k^{5/3})$ . One possibility for recovering  $\mathcal{O}(n_k)$  element assembly is to replace  $R_i$  with a matrix  $\tilde{R}_i$ , in which the values of hanging nodes only depend on the nearest independent nodes, as illustrated in fig. 5.3. The convergence plots in figs. 5.1 and 5.2 were generated on the mesh  $\mathcal{T}_{xy}$ , which is the default mesh for our test problems, using these sparse-approximation restriction matrices  $\{\tilde{R}_i\}$ : in this case, their use does not affect the convergence. If we use the mesh  $\mathcal{T}_{xyz}$  (fig. 4.1c), which is generated by standard isotropic octree-based refinement, the errors incurred by these sparse restriction operators increase as  $\phi$  in-






FIG. 5.3: Illustration of different constructions of the low-order matrix  $\tilde{F}$  at nonconforming interfaces.  $\tilde{F}$  can be thought of as the matrix for a finite element space whose functions are piecewise linear between the nodes of the original high-order finite element space. Because these nodes do not align at nonconforming interfaces (see fig. 3.2), the functions in this low-order space have discontinuities. The nature of these discontinuities depends on how the nodal values of the smaller element, shown in green, are interpolated from the global nodes, which coincide with the nodes of the larger element, shown in red. (a) If the same restriction operator  $R_i$  is used as for the original high-order finite element space, dependent hanging nodes are interpolated as though the independent nodes represent high-order polynomials. Each hanging node depends on every independent node along the nonconforming edge/face. (b) If the approximate restriction operator  $\tilde{R}_i$  is used, dependent hanging nodes are interpolated linearly (or bilinearly for two-dimensional faces) from the closest independent nodes. This results in sparser matrices and faster global matrix construction.




FIG. 5.4: Relative residual  $\|r_j\|/\|r_0\|$  versus number of Krylov iterations for the test problems described in section 4.3 for the mesh  $\mathcal{T}_{xyz}$  (fig. 4.1c), which has nonconforming faces in all directions. We compare two possibilities of handling nonconforming interfaces in the construction of the low-order preconditioner  $\tilde{F}$  at different element aspect ratios. The first possibility is to use  $\{R_i\}$ , the same hanging node restriction matrices as in the high-order discretization. An alternative is to use  $\{\tilde{R}_i\}$ , a piecewise linear approximation of  $\{R_i\}$ .

creases. Using the true  $\{R_i\}$  for  $\mathcal{T}_{xyz}$  improves the convergence, but we still observe slight  $\phi$ -dependence (see fig. 5.4).

To understand this behavior, note that in essence we are comparing two approximations to the high-order element matrix  $R_i^T F_i R_i$ : one,  $R_i^T \tilde{F}_i R_i$ , where the element matrix is replaced by a low-order approximation, and another,  $\tilde{R}_i^T \tilde{F}_i \tilde{R}_i$ , where additionally the high-order interpolation of the nodal values at nonconforming faces is replaced by a low-order interpolation. The latter approximation only affects nonconforming faces; the fact that nonconforming interfaces are much larger in  $\mathcal{T}_{xyz}$  than




FIG. 5.5: Comparison of the generalized eigenvalues  $\lambda$  that satisfy  $F\mathbf{u} = \lambda\tilde{F}\mathbf{u}$  for two meshes,  $\mathcal{T}_{xy}$  and  $\mathcal{T}_{xyz}$ , and for  $\tilde{F}$  constructed with either the true restriction operator  $R_i$  or its sparse approximation  $\tilde{R}_i$ . (a) The generalized eigenvalues for the mesh  $\mathcal{T}_{xy}$  are almost the same whether  $R_i$  or  $\tilde{R}_i$  is used. (b) The generalized eigenvalues for the mesh  $\mathcal{T}_{xyz}$  are sensitive to the use of  $\tilde{R}_i$ .

in  $\mathcal{T}_{xy}$  for large  $\phi$  seems to explain the different convergence behavior for problems discretized on the two meshes.

To investigate the influence of nonconforming faces on the low-order preconditioner numerically, we consider the generalized eigenvalue equation  $F\mathbf{u} = \lambda\tilde{F}\mathbf{u}$  for  $\phi = 100$  on the meshes  $\mathcal{T}_{xy}$  and  $\mathcal{T}_{xyz}$ ; the results are shown in fig. 5.5. For eigenvectors  $\mathbf{u}$  with eigenvalues far from 1,  $\tilde{F}$  is a poor approximation of  $F$ . For the mesh  $\mathcal{T}_{xy}$ , the errors incurred by the sparse restriction matrices  $\{\tilde{R}_i\}$  due to the nonconforming interfaces are small. For  $\mathcal{T}_{xyz}$ , errors incurred by  $\{\tilde{R}_i\}$  are more significant, and inspection of the extremal eigenvalues shows that they are associated with vectors  $\mathbf{u}$  for which  $F\mathbf{u}$  is large only on nonconforming interfaces that are normal to the  $z$ -axis.

**5.2.3. Residual computation for smoother on the finest mesh.** When the matrix  $\tilde{F}$  used to generate a multigrid hierarchy differs from the true matrix  $F$ , one has two possibilities for the residual computation in the smoothing step on the finest mesh. Namely, one can use the high-order discretized operator, i.e.,  $r = b - Fx$ , or its low-order approximation, i.e.,  $r = b - \tilde{F}x$ . The convergence in fig. 5.6a, which used  $\tilde{F}$  to define the AMG residuals, should be compared to fig. 5.1b, which uses  $F$  to define residuals and converges faster. Thus, in the following, we use the high-order operator for the residual computation in the smoother on the finest mesh.

**5.3. Convergence for different orders  $k$ .** In fig. 5.6b we show the convergence of our solver (GMRES, preconditioned by column-preserving SA-AMG with a damped block-Jacobi/IC(0) smoother) for polynomial orders  $k=3, 6$ , and  $9$  on a mesh with  $\phi = 100$ . As can be seen, the iteration number is independent of the polynomial order. To illustrate that the preconditioner is also  $k$ -independent with respect to computational work, note that the operator complexities of the AMG hierarchies with respect to  $\tilde{F}$  (i.e., the sum of nonzero matrix entries in all operators in the AMG hierarchy divided by the number of nonzero entries in  $\tilde{F}$ ) reported by PETSc are 1.22, 1.26, and 1.27 for  $k = 3, 6$  and  $9$ , respectively, and that the average numbers of nonzeros per row in  $\tilde{F}$  are 86.83, 82.33, and 82.36. This demonstrates that the cost to construct  $\tilde{F}$  and the coarse hierarchy are proportional to the problem size,

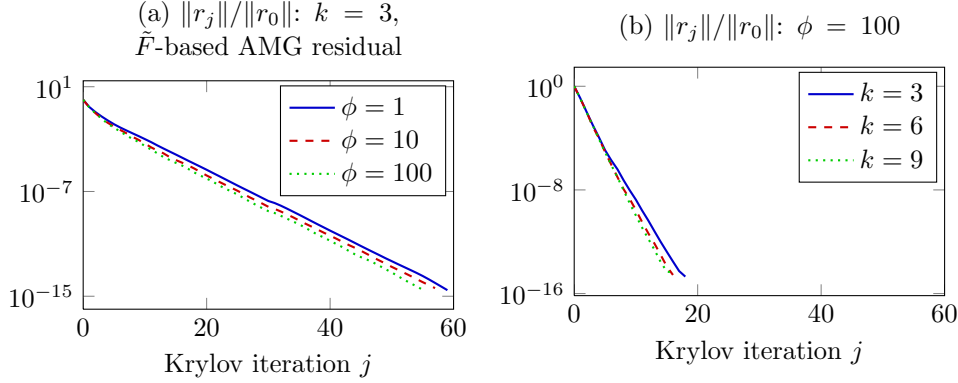


FIG. 5.6: Relative residual  $\|r_j\|/\|r_0\|$  versus number of Krylov iterations for the test problems described in section 4.3 on mesh  $\mathcal{T}_{xy}$  using SA-AMG smoothed by block-Jacobi/IC(0), (a) when the residual in the smoother on the finest level of the AMG hierarchy is computed with  $\tilde{F}$  instead of  $F$ , and (b) for varying orders of approximation.

but independent of  $k$ . The same can be said of the cost of computing the incomplete factorization smoother on each level as we do not allow fill-in. Thus, the overall computational complexity for preconditioning  $F$  is independent of  $k$ .

**6. Preconditioning the Schur complement of the (1,1)-block.** In this section, we analyze how different approximations of the Schur complement  $S$  of the (1,1)-block and different choices of the basis for the pressure space  $\mathcal{M}$  affect the convergence rate for Stokes linear problems preconditioned as in eq. (4.2). For all our tests, the preconditioner for the (1,1)-block  $F$  is a single multigrid V-cycle with the parameter choices described above in section 5. We test the effectiveness of our preconditioner on Stokes problems whose setup is discussed in section 4.3. As discussed in section 3.2, the discrete inf-sup constant for the  $\mathbb{Q}_k \times \mathbb{P}_{k-1}$  mixed element is  $\phi$ -dependent. In fig. 6.1a, we demonstrate this  $\phi$ -dependence numerically. Because of this instability we prefer the  $\mathbb{Q}_k \times \mathbb{Q}_{k-2}$  mixed element for anisotropic problems.

For constant viscosity  $\mu > 0$ ,  $S$  is known to be spectrally equivalent to the scaled pressure mass matrix  $-\mu^{-1}M$ . Because of this equivalence, a common choice is to approximate  $S$  by  $-\tilde{M}(\mu^{-1})$ , which is a diagonally-lumped approximation to the  $\mu^{-1}$ -weighted mass matrix. This Schur complement approximation does not take into account the anisotropic part of the 4th-order tensor  $\boldsymbol{\mu}'(\mathbf{u})$  defined in (2.7). As an alternative to the weighted mass matrix, the least-squares commutator, also known as the BFBt preconditioner, has proven to be a good Schur complement approximation, in particular in the presence of strongly varying coefficients [19, 43].

We have implemented both of these Schur complement preconditioners, and have found  $-\tilde{M}(|\mu|^{-1})$  to be the most efficient for the problems targeted in this paper. One reason for this good performance is that we choose a basis for the discontinuous finite element pressure space that nearly diagonalizes the mass matrix  $M$  for the pressure space. As a consequence, the effect of replacing  $-\tilde{M}(|\mu|^{-1})$  with its mass-lumped counterpart  $-\tilde{M}(|\mu|^{-1})$  is minimal. Such a basis for  $\mathbb{Q}_{k-2}(\hat{K})$  is given by a Lagrange basis for the tensor-product Gauss nodes. For a mapped element  $K_i$ , the mass matrix remains nearly diagonal provided the mapping is moderately nonaffine. To illustrate the effect of the choice of the basis, in fig. 6.1b we compare the convergence when




FIG. 6.1: Relative residual versus number of Krylov iterations for a Stokes test problem, posed on the same domain as the test problems in section 4.3, but with a uniformly refined mesh. The convergence obtained with different discretizations or bases for the pressure space are compared. (a) Comparison between the pressure spaces  $\mathbb{Q}_{k-2}$  and  $\mathbb{P}_{k-1}$ . (b) Comparison between Gauss points and Gauss-Lobatto points for nodal bases of the pressure space  $\mathbb{Q}_{k-2}$ .

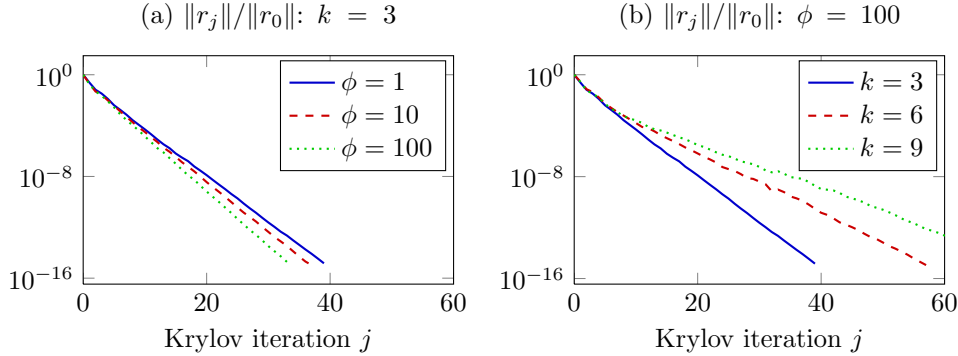


FIG. 6.2: Relative residual versus number of Krylov iterations for Stokes test problems. (a) The convergence for different aspect ratios  $\phi$  is compared. (b) The convergence for different polynomial orders  $k$  is compared.

using this Lagrange basis for the Gauss nodes to the convergence when using the Lagrange basis for Gauss-Lobatto nodes, which is more commonly used as a basis for tensor-polynomial finite elements.

In fig. 6.2a, we observe that the preconditioner for the Stokes system  $A$  results in convergence independent of the aspect ratio  $\phi$ ; however, as shown in fig. 6.2b, we find a dependence of the convergence on  $k$ . Note that this differs from our findings obtained for the (1,1)-block shown in fig. 5.6b, which shows independence of  $k$  for solving systems with  $F$ . The  $\mathcal{O}(k^{-1})$  decay of the lower bound for the discrete inf-sup constant for  $\mathbb{Q}_k \times \mathbb{Q}_{k-2}$  suggests that this  $k$ -dependence cannot be avoided.

**7. Nonlinear ice stream problems with smooth and rough beds.** To test the nonlinear solver for eq. (2.5), we adapt a model problem from [13]. As in




FIG. 7.1: Ice stream model problem on a  $400 \times 400$  km periodic domain of 1 km thick ice with a  $0.5^\circ$  slope to the right. Shown in (a) is the Robin coefficient field  $\beta$  beneath the ice and in (b) the magnitude of the velocity  $\mathbf{u}$  at the top surface. A fast flowing ice stream develops due to the shear thinning rheology.

section 4.3, the domain is a cutout of an infinite slab that is periodic in the horizontal dimensions, but the pitch of the domain relative to the direction of gravity is  $\theta = 0.5^\circ$ , so that a flow is induced. The Robin coefficient field  $\beta$  is shown in fig. 7.1a. Although  $\beta$  varies smoothly, the nonlinearity of the rheology causes the velocity  $\mathbf{u}$  to develop a narrow region of fast flow similar to an ice stream, as shown in fig. 7.1b. The constants in the constitutive relationship eq. (2.2) are  $n = 3$  and  $B(T) \equiv 2.15 \times 10^5 \text{ Pa a}^{1/3}$ , which equals  $A^{-1/3}$ , where  $A = 10^{-16} \text{ Pa}^{-3} \text{ a}^{-1}$ , which is taken from [49]. Because the top surface of the periodic domain is flat, we can convert the body force due to gravity into a constant tangential surface force with magnitude  $\rho g \sin \theta$ : we use an ice density of  $\rho = 910 \text{ kg m}^{-3}$ , and the acceleration due to gravity is  $g = 9.81 \text{ m sec}^{-2}$ . Note that with this change in the forcing, the pressure  $p$  is now interpreted as the variation from hydrostatic pressure. We use  $\varepsilon = 1 \times 10^{-6} \text{ a}^{-2}$ , which has a negligible effect for stresses of  $10^5 \text{ Pa}$  or greater, following the recommendation in [31, Chapter 2]. The periodic domain is  $400 \text{ km} \times 400 \text{ km} \times 1 \text{ km}$ , and we again use  $\mathcal{T}_{xy}$  as our mesh so that the elements are stretched to  $\phi = 100$ .

In fig. 7.2, we show the convergence behavior of the inexact Newton method for  $k = 3, 4$ , and 5, and compare with the convergence of an inexact Picard method for  $k = 3$ . As can be seen, the inexact Newton method converges faster than Picard’s method, both in terms of the number of nonlinear iterations and in terms of the total number of Krylov iterations. As each nonlinear iteration requires a preconditioner setup (or update), the superiority of Newton’s method compared to the Picard method is even more pronounced if we consider time-to-solution rather than the number of Krylov iterations. Although not shown in fig. 7.2, we did test Picard’s method with tighter tolerances on the linear solves, but found that this did not improve the convergence rate in terms of total Krylov iterations.

We next test our method on the same problem, but with the Robin coefficient field reduced to 1% of the previous field, and with rough bed geometries instead of the flat slab used in the previous test. These modifications make the problem more realistic, but more challenging to solve. We generate bed topographies using random coefficients in a truncated Fourier series, as in eq. (4.3). By changing the exponent  $\gamma$ , which controls the decay of the Fourier coefficients, we are able to control how rough the generated topography is. In figs. 7.3a and 7.3b, we show two topographies, generated with ten Fourier modes and  $\gamma = 1.5$  and  $\gamma = 1.0$ , respectively. In fig. 7.3c we




FIG. 7.2: Convergence of the inexact Newton solver for the flat bed topography ice stream model problem, for different polynomial orders  $k$ . Diamonds correspond to the nonlinear residual at the start of a Newton step. The solid lines show the linear residuals of the Krylov iterations. The dots show where the nonlinear residuals are evaluated and the preconditioners are recomputed. The Krylov method is FGMRES(30). The smoother for the (1,1)-block AMG preconditioner is two iterations of CG, preconditioned by IC(0) on each processor block. As can be observed, the linear and the nonlinear convergence rates coincide close to the solution, as desired for an inexact Newton-Krylov method. Far from the solution, where the linear and nonlinear residuals are very different, the inexactness in the method avoids over-solving of the linear systems, and is thus efficient in terms of Krylov iterations [18]. The bracketed numbers in the legend indicate the total number of nonlinear iterations.

show the convergence behavior of our method on domains with these bed topographies.

In these nonlinear problems, the incomplete factorization of the (1,1)-block approximation  $\tilde{F}$  sometimes encounters zero or negative pivots on the diagonal, which can lead to poor convergence or can cause the solver to fail. Zero or negative pivots occur more likely when there are regions of rough topography and sharp solution gradients. We have considered several strategies to make the preconditioner robust in these cases. One remedy to avoid these bad pivots is to increase the sophistication of the incomplete factorization, for instance by increasing the level of fill-in (e.g., use IC(1) or IC(2) instead of IC(0)), or by using a drop-tolerance-based approach. However, these approaches incur additional setup and storage costs, and it is difficult to anticipate a priori what amount of fill (or what drop tolerance) is needed to avoid bad pivots.

PETSc implements strategies for modification of the matrix being factorized to avoid bad pivots. The default is to use Manteuffel’s shifting strategy [42], in which  $\tilde{F}$  is replaced by  $\tilde{F} + \alpha I$ . Our experience is that when the magnitudes of the entries on the diagonal of  $\tilde{F}$  are highly variable, which can be due to variable coefficients or variable element sizes, this shifting strategy is detrimental to the effectiveness of the incomplete factorization as a smoother. In particular, it can shift the eigenvalues of the high-frequency modes away from the region that is optimally damped by a Chebyshev polynomial smoother, leading to stagnation in the convergence. We find that shifting

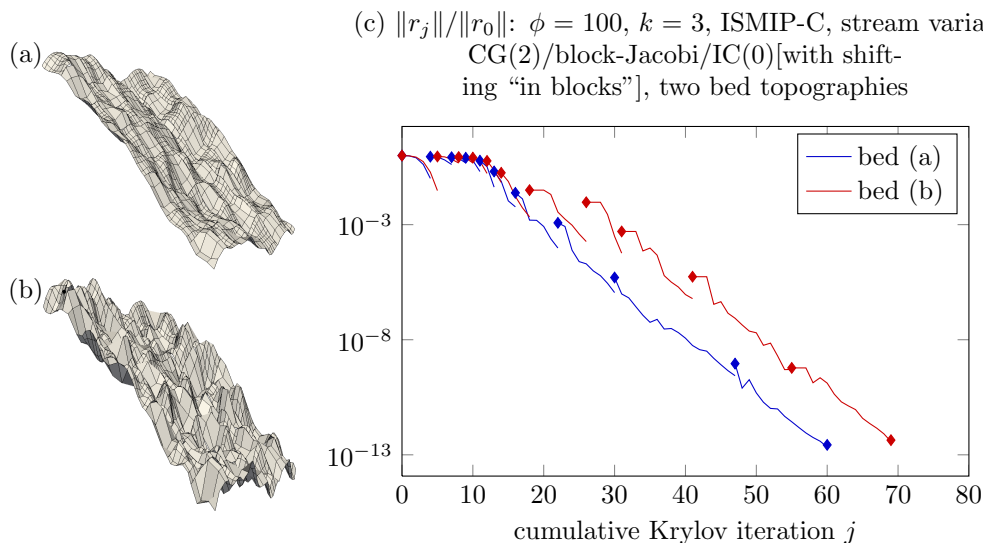


FIG. 7.3: (a), (b) Bed topographies with different roughness for the ice stream model problem. The vertical scale is exaggerated by a factor of 100. The maximum variation in the bed topography in (b) is about half of the ice thickness. (c) Convergence of the inexact Newton solver for the ice stream model problem as described in fig. 7.1, but with 1% of the basal friction  $\beta$  and the rough bed topographies.

“in blocks”—increasing pivots to a lower bound when they become too small—results more consistently in factorizations that are appropriate for use as smoothers. The corresponding convergence behavior for problems with polynomial order  $k = 3$  is shown in fig. 7.3. Note that while compared to fig. 7.2, the efficiency in terms of overall Krylov is similar, the inexact Newton method requires more iterations before it reaches the asymptotic convergence regime. This is likely due to a combination of a greater degree of nonlinearity in the problem caused by the more complex geometry and the larger variations in the Robin coefficient  $\beta$ .

**8. Antarctic ice sheet problem.** We now demonstrate the performance of our inexact Newton-Krylov solver for the simulation of the dynamics of the Antarctic ice sheet. Below, we detail how satellite and radar data are used to define the computational domain and describe the mesh generation. In section 8.2, we study the performance of our solver for this problem. Finally, in section 8.3, we compare to results presented in [33], which are obtained with a similar solver that uses a SSOR smoother in the preconditioner instead of an incomplete factorization smoother.

**8.1. Problem Description.** For the following simulations, we define the ice density to be  $\rho = 917 \text{ kg/m}^3$ , the pre-exponential in Glen’s power law (2.2) to be  $B(T) = 4.1 \times 10^5 \text{ Pa a}^{1/3}$ , and the regularizing constant that prevents infinite effective viscosity to be  $\epsilon = 9.95 \times 10^{-6} \text{ a}^{-2}$ . The assumed sliding coefficient field  $\beta$  is computed by first taking the ratio of the driving stress due to gravity and the observed surface velocity, and then imposing maximum and minimum values. This results in  $\beta$  being  $\sim 0.3 \text{ MPa a km}^{-1}$  over most of the ice sheet and becomes almost zero (with a minimum value of  $\sim 10^{-14} \text{ MPa a km}^{-1}$ ) in ice streams.

The body force due to gravity is  $-\rho g \mathbf{z}$ , where  $\mathbf{z}$  is a unit vector that is normal to



the reference ellipsoid. Due to the complexity of the ice sheet geometry, this forcing cannot be exactly represented by the discrete function spaces we use. Consequently, we find that when we solve eq. (2.5) with the right-hand side computed by directly discretizing the body force due to gravity, the velocity in our solution exhibits grid-scale oscillations. To ameliorate this, we set the body force to be  $\mathbf{b} = -\rho g \mathbf{z} - \text{grad } \hat{p}$ , where  $\hat{p}$  is a hydrostatic-like pressure solving

$$\begin{aligned} -\Delta \hat{p} &= 0 & \mathbf{x} \in \Omega, \\ \hat{p} &= 0 & \mathbf{x} \in \Gamma_N, \\ \nabla \hat{p} \cdot \mathbf{n} &= -\rho g \mathbf{z} \cdot \mathbf{n} & \mathbf{x} \in \Gamma_D, \\ \nabla \hat{p} \cdot \mathbf{n} &= -\rho g \mathbf{z} \cdot \mathbf{n} & \mathbf{x} \in \Gamma_R. \end{aligned}$$

Note that  $\hat{p}$  satisfies homogeneous Dirichlet boundary conditions where the velocity field has stress-free boundary conditions, and  $\hat{p}$  has Neumann boundary conditions where the normal velocity component satisfies a Dirichlet condition. The pressure  $p$  that is solved for is then a variation from  $\hat{p}$ . Note that the computation of  $\hat{p}$  is an upfront computation and is not included in any of the performance results that follow.

The geometric description of the ice sheet is constructed from the ALBMAP dataset [40]. Elevation values in the ALBMAP dataset are given relative to the EIGEN-GL04C geoid [21]: we convert these values to elevations relative to the WGS84 ellipsoid [1] using the software library GeographicLib [38], and then map the resulting (latitude, longitude, elevation) geodetic coordinates into Cartesian coordinates. While this results in a usable geometry for the ice sheet, the resulting geometry for the ice shelves, the extension of the ice sheet onto the surface of the ocean, is far from hydrostatic equilibrium, resulting in flow velocities several orders of magnitude too fast. We were unable to correct this behavior using local geometry adjustments. While we believe that a good geometry can be obtained, this was not further pursued for these numerical studies and we have limited the domain to the grounded ice sheet.

From the ice thickness data, given on a latitude-longitude grid, we obtain a polygon describing the lateral boundaries of the ice sheet. We create a quadrilateral mesh from this polygon by first using the triangular mesh generator Triangle [55], and then splitting the triangles into quadrilaterals, to which we apply mesh smoothing to improve the element quality. This resulting coarse quadrilateral mesh (fig. 8.1a) contains  $\sim 27,000$  elements. Using the quadtree-based refinement for the horizontal directions within our `p4est` extension for anisotropic domains, we refine this mesh to construct the footprints for the columns of our final hexahedral mesh. We use several refinement criteria: we require that the elements of our mesh have a footprint smaller than  $(2.5 \text{ km})^2$  at this grounding line; we refine any column whose thickness varies by more than a factor of 1.5; in keeping with other ice sheet models, we keep the aspect ratio  $\phi$  of the elements below 25. Once we have constructed a mesh that satisfies these constraints (fig. 8.1b), we use uniform refinement (i.e., replace each hexahedron by eight children) from this mesh when we perform scaling studies.

The Antarctic ice sheet contains some very thin regions which, because we constrain the aspect ratio of our elements, would require a large number of elements if we modeled the true thickness of the ice sheet. To control the mesh size, we employ an artificial minimum thickness of 200 m, enforced by modifying the bedrock topography.

## 8.2. Solver performance and scalability.

*Nonlinear convergence for different orders  $k$ .* We first test our solver for different polynomial orders on a mesh obtained by one level of uniform mesh refinement. For




FIG. 8.1: Antarctic ice sheet problem: (a) The coarse quadrilateral mesh description of the Antarctic ice sheet. (b) An oblique view of the refined hexahedral mesh, obtained using the “*p6est*” extension of the *p4est* library for anisotropic AMR. The refinement near the center is the result of mesh refinement to reduce the aspect ratio of elements in the mesh: this type of refinement is not possible using purely octree-based refinement.




FIG. 8.2: Antarctic ice sheet problem: convergence of the incomplete Newton-Krylov solver for different polynomial element orders  $k$ . We enforce a minimal ice thickness of 200 m, and the mesh is distributed over 1024 processes. The bracketed numbers are the total number of nonlinear iterations.

smoothing in the (1,1)-block multigrid preconditioner, we use 2 GMRES iterations of our block-Jacobi/IC(0) smoother. We discretize eq. (2.5) using  $\mathbb{Q}_k \times \mathbb{Q}_{k-2}$  elements for  $k = 3, 4$ , and 5, resulting in problems with 51M, 121M, and 238M degrees of freedom. The experiments were conducted on TACC’s Stampede supercomputer, with each MPI process assigned to one Sandy Bridge Xeon core. Each discretization is distributed across 1024 MPI processes. The convergence of our inexact Newton-Krylov solver is shown in fig. 8.2. Note that the convergence is similar to the results for the rough bed topography model problem shown in fig. 7.3c. The increased number of overall Krylov iterations (by about a factor of 2) is likely due to the more complex geometry and boundary conditions.

*Parallel scalability.* We next test the parallel scalability of our solver on the  $\mathbb{Q}_3 \times \mathbb{Q}_1$  discretization. We use three different levels of uniform mesh refinement,  $\ell = 0, 1$ , and 2. In conducting our study, we found that our GMRES(2)/block-

TABLE 8.1: *Antarctic ice sheet problem: scaling of the full nonlinear solver for  $k = 3$  for the geometry with minimal ice thickness of 200 m, using GMRES(10)/block-Jacobi/IC(0) as a smoother in the (1,1)-block multigrid preconditioner. For different numbers of processes  $P$ , we report the total time (in seconds) to solve the nonlinear problem to a relative tolerance of  $10^{-12}$  in the  $\ell^2$ -norm from a zero initial guess, the parallel efficiency (eff.) as well as the number of Newton iterations ( $\#N$ ) and overall Krylov iterations ( $\#K$ ) performed during the solution. We also report the average time and efficiency for a single Stokes operator application (Op), a preconditioner application (PC), and the total preconditioner setup time (setup). Factorization of the coarsest operator (which has always less than 50 unknowns) is included in the setup times, and applying the inverse of the factored coarse operator is included in preconditioner application times. We report efficiency in reference to the smallest problem size on the smallest number of processors.*

$P$	Solve	eff.	$\#N$	$\#K$	Op	eff.	PC	eff.	Setup	eff.
$\ell = 0$ uniform refinement, $N_{\text{dof}} = 7\text{M}$										
128	67.7	1.00	7	66	0.0275	1.00	0.981	1.00	7.361	1.00
256	36.9	0.91	7	67	0.0146	0.94	0.577	0.93	4.255	0.86
512	20.3	0.83	7	65	0.0080	0.86	0.299	0.82	2.722	0.67
$\ell = 1$ uniform refinement, $N_{\text{dof}} = 51\text{M}$										
1,024	78.0	0.86	11	75	0.0276	0.99	0.987	0.99	8.657	0.85
2,048	44.2	0.76	10	75	0.0152	0.90	0.561	0.87	5.810	0.63
4,096	30.2	0.56	10	75	0.0087	0.79	0.383	0.64	5.406	0.34
$\ell = 2$ uniform refinement, $N_{\text{dof}} = 383\text{M}$										
8,192	108.0	0.62	13	91	0.0295	0.93	1.13	0.87	15.58	0.47
16,384	74.7	0.45	10	89	0.0168	0.82	0.80	0.61	17.22	0.21

Jacobi/IC(0) smoother for the (1,1)-block multigrid preconditioner was sometimes insufficient: because of the non-smooth coefficients obtained for some Newton steps, the prolongation of coarse solutions onto fine grids would introduce errors that could not be sufficiently damped by just a couple of smoother iterations. For this reason, we employ a stronger GMRES(10)/block-Jacobi/IC(0) smoother in this scaling study. We report the results in table 8.1. The table includes timings for each of the main components of the nonlinear solver: the matrix-vector product, the preconditioner application, and the preconditioner reconstruction.

The time to apply the preconditioner is almost entirely spent in applying the AMG V-cycle to the (1,1)-block of the Stokes operator. This step also requires more communication than the matrix-vector product, both to project and restrict vectors in the hierarchy and to apply the smoothers. Note that GMRES requires global reductions to compute the required inner products, which would not be the case for a stationary smoother (e.g. Chebyshev).

It is well known that the setup phase in parallel implementations of algebraic multigrid requires significant communication to properly aggregate degrees of freedom across processor boundaries, compute prolongation matrices and repartition coarse matrices. As a consequence, the setup is often the least scalable component of the solver [8, 12], which can also be seen in table 8.1. Here, the reported times include both the prolongator constructions in the initial setup and the Galerkin projections to reconstruct the coarse matrices during subsequent Newton steps, so the increase in Newton iterations for the larger problems affects the reported efficiency. In all cases, the contribution of the initial prolongator setup is roughly equivalent to the cost of

one subsequent preconditioner reconstruction.

Note that though the total number of Krylov iterations increases by  $\sim 37\%$  from the smallest to the largest problem, the number of outer inexact Newton steps almost doubles from 7 to 13. Recent work [7] on solving the hydrostatic approximation to the Stokes equations for ice sheet dynamics has demonstrated the effectiveness of grid continuation in obtaining an initial guess that is near the region of asymptotic convergence of Newton’s method. The hierarchical mesh refinement we use lends itself naturally to grid continuation, so this approach could improve the efficiency of our nonlinear solver.

**8.3. Comparison with performance of SSOR-based smoothing.** In [33], we have used a similar nonlinear and linear solver framework as the one developed in this paper to infer the Robin coefficient field  $\beta$  in the Stokes boundary value problem such that the velocity fields of the solution closely matches satellite observations of the Antarctic ice sheet’s surface (see fig. 8.3). This inference uses methods of adjoint-based PDE-constrained optimization to find an optimal  $\beta$ , which requires the solution of the Stokes problem described in this work, as well as linearized adjoint problems, whose operators are similar to the linearized Stokes operator in (2.6).

The difference between the solver used in [33] and the one presented here lies in the preconditioning for the (1,1)-block. In [33], a Chebyshev(2)/block-Jacobi/SSOR smoother combined with the standard SA-AMG scheme (not the column-preserving SA-AMG presented in section 5.1) is used. In table 8.2, we reproduce some of the algorithmic scalings from [33], and compare them to scalings obtained for the same problem with the solvers developed here: using a Chebyshev(2)/block-Jacobi/IC(0) smoother and column-preserving SA-AMG. We present only algorithmic scalability because the results in [33] were computed on Oak Ridge National Laboratory’s Titan supercomputer, while the new scalings were computed on TACC’s Stampede supercomputer. As a consequence, the timings are not directly comparable. We note, however, that the cost of applying an IC(0) preconditioner, in terms of memory movement and floating point operations, is almost identical to the cost of applying an SSOR preconditioner, so the only significant additional computation in our new scaling results is in the computation of the IC(0) factorization, which is on par with a single application of the IC(0) preconditioner. Thus, the algorithmic speedup shown in table 8.2 is representative of the runtime speedup we would find, had these scaling results been produced on the same architecture.

We note that the total number of iterations required to solve these problems is lower than for the problems solved earlier in fig. 8.2, even though both are posed on the same domain. This is because the aspect ratio of the elements is smaller, with all element aspect ratios  $\phi$  less than 10, and because the basal friction is larger, reaching  $\sim 3 \text{ MPa a km}^{-1}$  in some areas (versus the  $\sim 0.3 \text{ MPa a km}^{-1}$  used in fig. 8.2), and having a minimum value in the ice streams that is three orders of magnitude larger than the minimum used in fig. 8.2.

**9. Conclusions.** Several issues related to high-order, adaptive mesh discretizations and solvers for the simulation of nonlinear Stokes flow in three-dimensional anisotropic domains are addressed in this work. Our main target problem has been the nonlinear Stokes boundary value problem arising in ice sheet dynamics. We demonstrate an extension to the `p4est` library for adaptive mesh refinement of anisotropic domains that combines quadtree-based refinement in the horizontal directions with columns of elements to achieve a flexible approach to mesh refinement, with local control over the aspect ratio of elements in the mesh. We demonstrate that high-order




FIG. 8.3: *Antarctic ice sheet problem: magnitude of the surface velocity field optimized to match satellite observations. The inversion procedure to infer the basal coefficient is described in [33].*

TABLE 8.2: *Antarctic ice sheet problem: comparison of algorithmic scalability of the solver used in [33] (which uses SSOR and standard SA-AMG within the (1,1)-block preconditioner) versus the one presented here (which uses IC(0) and column-preserving SA-AMG). Shown are the number of overall degrees of freedom  $N_{\text{dof}}$ , the number of processes  $P$ , the number of Newton iterations  $\#N$  and the number of cumulative Krylov iterations  $\#K$ . The data on the left is reproduced from [33].*

		SSOR, standard SA-AMG		IC(0), column-preserving SA-AMG	
$N_{\text{dof}}$	$P$	$\#N$	$\#K$	$\#N$	$\#K$
38M	1,024	8	147	7	85
268M	8,192	9	243	8	98

finite elements discretized on these meshes are well-approximated by low-order approximations, which can be used for preconditioning. We present an efficient solver for the linearized Stokes equations, with particular emphasis on the design of algebraic multi-grid solvers for high-order discretizations, anisotropic domains, and hanging nodes. Using incomplete factorization-based smoothing for the (1,1)-block yields efficient and fast convergence. When using this type of smoothing, we demonstrate that column-preserving SA-AMG, as implemented by our DofColumns plugin for PETSc’s GAMG preconditioner, significantly improves over standard SA-AMG in its effectiveness, especially in the presence of weak boundary conditions. The numerical experiments on our discretization of the Antarctic ice sheet show that, up to a point, the incomplete factorization process can be made robust against bad pivots that occur in problems with variable coefficients and rough topography by using only local shifting and by combining the incomplete factorization with a non-stationary Krylov method as a smoother.

**10. Acknowledgments.** We would like to thank the three referees who reviewed this work, whose reviews greatly improved the final version of this work. In particular, we would like to thank the two initial referees, Ray Tuminaro and one who remains anonymous: the column-preserving SA-AMG that we have developed in the DofColumns plugin was not included in the initial submission of this work, but was developed in response to their insightful comments and suggestions.

**Appendix A. Well-posedness of eq. (2.5).** The main difference between eq. (2.5) and the variational form in [35] is the boundary integral  $\int_{\Gamma_R} \beta T_{\parallel} \mathbf{v} \cdot T_{\parallel} \mathbf{u} \, ds$ :

for  $\mathbf{u} \in \mathcal{W}^{1,r}(\Omega)$ , this form is not meaningful because the trace of  $\mathbf{u}$  is not necessarily in  $\mathcal{L}^2(\Gamma_R)$ . We compensate for this by bringing the boundary integral into the definition of  $\mathcal{V}$ . Let  $[\mathcal{C}^\infty(\Omega)]_0^3$  be the space of smooth vector-valued functions in  $\Omega$  that satisfy the homogeneous Dirichlet part of the boundary conditions and let  $r = 1 + \frac{1}{n}$ . Under the above assumptions, the functional

$$\mathcal{J}(\mathbf{u}) = \left\{ \int_{\Omega} |\nabla \mathbf{u}|^r d\mathbf{x} \right\}^{1/r} + \left\{ \int_{\Gamma_R} \beta |\mathbf{T}_{\parallel} \mathbf{u}|^2 ds \right\}^{1/2} \quad (\text{A.1})$$

defines a seminorm on  $[\mathcal{C}^\infty(\Omega)]_0^3$ . We assume a problem setup in which  $\mathcal{J}(\mathbf{u})$  also defines a norm, which amounts to requiring that if  $\mathbf{u} \in [\mathcal{C}^\infty(\Omega)]_0^3$  is a rigid-body motion that satisfies the Dirichlet conditions, then  $\int_{\Gamma_R} \beta |\mathbf{T}_{\parallel} \mathbf{u}|^2 ds > 0$ . We define  $\mathcal{V}$  to be the closure of  $[\mathcal{C}^\infty(\Omega)]_0^3$  in this norm.

Following the same steps as in [35], one can show that the minimization problem

$$\min_{\mathbf{u}} \frac{2n}{1+n} \int_{\Omega} B(T)(\mathbf{D}_{\Pi}(\mathbf{u}) + \varepsilon)^{\frac{1+n}{2n}} d\mathbf{x} + \frac{1}{2} \int_{\Gamma_R} \mathbf{T}_{\parallel} \mathbf{u} \cdot \mathbf{T}_{\parallel} \mathbf{u} ds - \mathbf{f}(\mathbf{u}) \quad (\text{A.2})$$

is well-posed in  $\mathcal{V}_{\text{div}}$ , the subspace of  $\mathcal{V}$  containing only the divergence-free functions. To prove that eq. (2.5) is well-posed, we additionally need to choose a space  $\mathcal{M}$  for which we can prove the inf-sup condition

$$\inf_{q \in \mathcal{M}} \sup_{\mathbf{u} \in \mathcal{V}} \frac{\int_{\Omega} q \nabla \cdot \mathbf{u} d\mathbf{x}}{\|q\|_{\mathcal{M}} \|\mathbf{u}\|_{\mathcal{V}}} \geq \gamma > 0. \quad (\text{A.3})$$

This inequality still holds for  $\mathcal{M} = \mathcal{L}^{r'}(\Omega)$ . To see this, consider the subspace  $\tilde{\mathcal{V}} = \{\mathbf{u} \in \mathcal{V} : \mathbf{u}|_{\Gamma_R} = 0\}$ : for  $\mathbf{u} \in \tilde{\mathcal{V}}$ ,  $\|\mathbf{u}\|_{\mathcal{V}} = \|\mathbf{u}\|_{[\mathcal{W}^{1,r}(\Omega)]^3}$ , and so

$$\inf_{q \in \mathcal{L}^{r'}(\Omega)} \sup_{\mathbf{u} \in \tilde{\mathcal{V}}} \frac{\int_{\Omega} q \nabla \cdot \mathbf{u} d\mathbf{x}}{\|q\|_{\mathcal{M}} \|\mathbf{u}\|_{\mathcal{V}}} \geq \inf_{q \in \mathcal{L}^{r'}(\Omega)} \sup_{\mathbf{u} \in \tilde{\mathcal{V}}} \frac{\int_{\Omega} q \nabla \cdot \mathbf{u} d\mathbf{x}}{\|q\|_{\mathcal{L}^{r'}(\Omega)} \|\mathbf{u}\|_{[\mathcal{W}^{1,r}(\Omega)]^3}}. \quad (\text{A.4})$$

As long as  $\Gamma_N \neq \emptyset$ , the term on the right is bounded from below as a particular case of the inf-sup condition in [35]. Thus, (2.5) is well posed and has a unique solution.

## REFERENCES

- [1] *Department of Defense World Geodetic System 1984: Its definition and relationships with local geodetic systems*, Tech. Rep. NGA TR8350.2, U.S. National Geospatial-Intelligence Agency, 1984. Revised in 2004.
- [2] M. F. ADAMS, *A parallel maximal independent set algorithm*, in Proceedings 5th copper mountain conference on iterative methods, 1998.
- [3] M. AINSWORTH AND P. COGGINS, *The stability of mixed hp-finite element methods for Stokes flow on high aspect ratio elements*, SIAM Journal on Numerical Analysis, 38 (2000), pp. 1721–1761.
- [4] S. BALAY, J. BROWN, K. BUSCHELMAN, V. ELJKHOUT, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNES, B. F. SMITH, AND H. ZHANG, *PETSc users manual*, Tech. Rep. ANL-95/11 - Revision 3.3, Argonne National Laboratory, 2012.
- [5] H. BLATTER, *Velocity and stress fields in grounded glaciers: a simple algorithm for including deviatoric stress gradients*, Journal of Glaciology, 41 (1995), pp. 333–344.
- [6] J. BROWN, *Efficient nonlinear solvers for nodal high-order finite elements in 3D*, Journal of Scientific Computing, 45 (2010), pp. 48–63.
- [7] J. BROWN, B. SMITH, AND A. AHMADIA, *Achieving textbook multigrid efficiency for hydrostatic ice sheet flow*, SIAM Journal on Scientific Computing, 35 (2013), pp. B359–B375.

- [8] C. BURSTEDDE, O. GHATTAS, G. STADLER, T. TU, AND L. C. WILCOX, *Parallel scalable adjoint-based adaptive solution for variable-viscosity Stokes flows*, Computer Methods in Applied Mechanics and Engineering, 198 (2009), pp. 1691–1700.
- [9] C. BURSTEDDE, L. C. WILCOX, AND O. GHATTAS, *p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees*, SIAM Journal on Scientific Computing, 33 (2011), pp. 1103–1133.
- [10] M. CAI, A. NONAKA, J. B. BELL, B. E. GRIFFITH, AND A. DONEV, *Efficient variable-coefficient finite-volume Stokes solvers*, Communications in Computational Physics, 16 (2014), pp. 1263–1297.
- [11] C. CANUTO, P. GERVASIO, AND A. QUARTERONI, *Finite-element preconditioning of G-NI spectral methods*, SIAM Journal on Scientific Computing, 31 (2010), pp. 4422–4451.
- [12] E. CHOW, R. D. FALGOUT, J. J. HU, R. S. TUMINARO, AND U. M. YANG, *A survey of parallelization techniques for multigrid solvers*, in Parallel Processing for Scientific Computing, M. A. Heroux, P. Raghavan, and H. D. Simon, eds., Society for Industrial and Applied Mathematics, Philadelphia, PA, 2006, pp. 179–201.
- [13] S. L. CORNFORD, D. F. MARTIN, D. T. GRAVES, D. F. RANKEN, A. M. LE BROCCQ, R. M. GLADSTONE, A. J. PAYNE, E. G. NG, AND W. H. LIPSCOMB, *Adaptive mesh, finite volume modeling of marine ice sheets*, Journal of Computational Physics, 232 (2012), pp. 529–549.
- [14] H. DAMANIK, J. HRON, A. OUAZZI, AND S. TUREK, *A monolithic fem-multigrid solver for non-isothermal incompressible flow on general meshes*, Journal of Computational Physics, 228 (2009), pp. 3869–3881.
- [15] M. DEVILLE AND E. MUND, *Finite-element preconditioning for pseudospectral solutions of elliptic problems*, SIAM Journal on Scientific and Statistical Computing, 11 (1990), pp. 311–342.
- [16] M. O. DEVILLE, P. F. FISCHER, AND E. H. MUND, *High-Order Methods for Incompressible Fluid Flow*, vol. 9 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, UK, 2002.
- [17] M. DURUFLE, P. GROB, AND P. JOLY, *Influence of Gauss and Gauss-Lobatto quadrature rules on the accuracy of a quadrilateral finite element method in the time domain*, Numerical methods for partial differential equations, 25 (2009), pp. 526–551.
- [18] S. C. EISENSTAT AND H. F. WALKER, *Choosing the forcing terms in an inexact Newton method*, SIAM Journal on Scientific Computing, 17 (1996), pp. 16–32.
- [19] H. ELMAN, V. HOWLE, J. SHADID, R. SHUTTLEWORTH, AND R. TUMINARO, *Block preconditioners based on approximate commutators*, SIAM Journal on Scientific Computing, 27 (2006), pp. 1651–1668.
- [20] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers with applications in incompressible fluid dynamics*, Oxford University Press, Oxford, 2005.
- [21] C. FÖRSTE, R. SCHMIDT, R. STUBENVOLL, F. FLECHTNER, U. MEYER, R. KÖNIG, H. NEUMAYER, R. BIANCALE, J.-M. LEMOINE, S. BRUINSMA, ET AL., *The Geoforschungszentrum Potsdam/groupe de Recherche de Geodesie Spatiale satellite-only and combined gravity field models: EIGEN-GL04S1 and EIGEN-GL04C*, Journal of Geodesy, 82 (2008), pp. 331–346.
- [22] M. FURUICHI, D. A. MAY, AND P. J. TACKLEY, *Development of a Stokes flow solver robust to large viscosity jumps using a Schur complement approach with mixed precision arithmetic*, Journal of Computational Physics, 230 (2011), pp. 8835–8851.
- [23] O. GAGLIARDINI AND T. ZWINGER, *The ISMIP-HOM benchmark experiments performed using the finite-element code Elmer*, The Cryosphere, 2 (2008), pp. 67–76.
- [24] M. W. GEE, J. J. HU, AND R. S. TUMINARO, *A new smoothed aggregation multigrid method for anisotropic problems*, Numer. Linear Algebra Appl, 16 (2009), pp. 19–37.
- [25] T. GEENEN, M. UR REHMAN, S. P. MACLACHLAN, G. SEGAL, C. VUIK, A. P. VAN DEN BERG, AND W. SPAKMAN, *Scalable robust solvers for unstructured FE geodynamic modeling applications: Solving the Stokes equation for models with large localized viscosity contrasts*, Geochemistry Geophysics Geosystems, 10 (2009), p. Q09002.
- [26] R. GLOWINSKI AND J. XU, *Numerical Methods for Non-Newtonian Fluids: Special Volume*, vol. 16 of Handbook of numerical analysis, North-Holland, 2011.
- [27] P. P. GRINEVICH AND M. A. OLSHANSKII, *An iterative method for the Stokes-type problem with variable viscosity*, SIAM Journal on Scientific Computing, 31 (2009), pp. 3959–3978.
- [28] V. HEUVELINE AND F. SCHIEWECK, *On the inf-sup condition for higher order mixed fem on meshes with hanging nodes*, ESAIM: Mathematical Modelling and Numerical Analysis, 41 (2007), pp. 1–20.
- [29] J. HEYS, T. MANTEUFFEL, S. MCCORMICK, AND L. OLSON, *Algebraic multigrid for higher-order finite elements*, Journal of computational Physics, 204 (2005), pp. 520–532.
- [30] R. C. A. HINDMARSH, *A numerical comparison of approximations to the Stokes equations used*



- in ice sheet and glacier modeling*, Journal of Geophysical Research, 109 (2004), p. F01012.
- [31] K. HUTTER, *Theoretical Glaciology*, Mathematical Approaches to Geophysics, D. Reidel Publishing Company, 1983.
  - [32] T. ISAAC, C. BURSTEDDE, AND O. GHATTAS, *Low-cost parallel algorithms for 2:1 octree balance*, in Proceedings of the 26th IEEE International Parallel & Distributed Processing Symposium, IEEE, 2012. <http://dx.doi.org/10.1109/IPDPS.2012.47>.
  - [33] T. ISAAC, N. PETRA, G. STADLER, AND O. GHATTAS, *Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet*, Journal of Computational Physics, 296 (2015), pp. 348–368.
  - [34] A. JANKA, *Smoothed aggregation multigrid for a Stokes problem*, Computing and Visualization in Science, 11 (2008), pp. 169–180.
  - [35] G. JOUVET AND J. RAPPAZ, *Analysis and finite element approximation of a nonlinear stationary Stokes problem arising in glaciology*, Advances in Numerical Analysis, 2011 (2012). ID 164581.
  - [36] Y. KANARSKA, A. SHCHEPETKIN, AND J. MCWILLIAMS, *Algorithm for non-hydrostatic dynamics in the regional oceanic modeling system*, Ocean Modelling, 18 (2007), pp. 143–174.
  - [37] E. KARER AND J. KRAUS, *Algebraic multigrid for finite element elasticity equations: Determination of nodal dependence via edge-matrices and two-level convergence*, International Journal for Numerical Methods in Engineering, 83 (2010), pp. 642–670.
  - [38] C. KARNEY, *Algorithms for geodesics*, Journal of Geodesy, 87 (2013), pp. 43–55.
  - [39] S. KIM, *Piecewise bilinear preconditioning of high-order finite element methods*, Electronic Transactions on Numerical Analysis, 26 (2007), pp. 228–242.
  - [40] A. M. LE BROCCQ, A. J. PAYNE, AND A. VIELI, *An improved Antarctic dataset for high resolution numerical ice sheet models (ALBMAP v1)*, Earth System Science Data, 2 (2010), pp. 247–260.
  - [41] W. LENG, L. JU, M. GUNZBURGER, S. PRICE, AND T. RINGLER, *A parallel high-order accurate finite element nonlinear Stokes ice-sheet model and benchmark experiments*, Journal of Geophysical Research, 117 (2012).
  - [42] T. A. MANTEUFFEL, *An incomplete factorization technique for positive definite linear systems*, Mathematics of computation, 34 (1980), pp. 473–497.
  - [43] D. A. MAY AND L. MORESI, *Preconditioned iterative methods for Stokes flow problems arising in computational geodynamics*, Physics of the Earth and Planetary Interiors, 171 (2008), pp. 33–47.
  - [44] D. MCKENZIE, *The generation and compaction of partially molten rock*, Journal of Petrology, 25 (1984), pp. 713–765.
  - [45] G. A. MEEHL, T. F. STOCKER, W. D. COLLINS, A. T. FRIEDLINGSTEIN, A. T. GAYE, J. M. GREGORY, A. KITOH, R. KNUTTI, J. M. MURPHY, A. NODA, S. C. B. RAPER, I. G. WATTERSON, A. J. WEAVER, AND Z.-C. ZHAO, *Global climate projections*, in Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, 2007, pp. 747–845.
  - [46] M. MORLIGHEM, H. SEROUSSI, E. LAROUR, AND E. RIGNOT, *Inversion of basal friction in antarctica using exact and incomplete adjoints of a higher-order model*, Journal of Geophysical Research: Earth Surface, 118 (2013), pp. 1746–1753.
  - [47] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer Verlag, Berlin, Heidelberg, New York, 1999.
  - [48] L. OLSON, *Algebraic multigrid preconditioning of high-order spectral elements for elliptic problems on a simplicial mesh*, SIAM Journal on Scientific Computing, 29 (2007), pp. 2189–2209.
  - [49] F. PATTYN, L. PERICHON, A. ASCHWANDEN, B. BREUER, B. DE SMEDT, O. GAGLIARDINI, G. H. GUDMUNDSSON, R. C. A. HINDMARSH, A. HUBBARD, J. V. JOHNSON, T. KLEINER, Y. KONOVALOV, C. MARTIN, A. J. PAYNE, D. POLLARD, S. PRICE, M. RUCKAMP, F. SAITO, O. SOUCEK, S. SUGIYAMA, AND T. ZWINGER, *Benchmark experiments for higher-order and full-Stokes ice sheet models (ISMIP-HOM)*, The Cryosphere, 2 (2008), pp. 95–108.
  - [50] N. PETRA, H. ZHU, G. STADLER, T. J. R. HUGHES, AND O. GHATTAS, *An inexact Gauss-Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model*, Journal of Glaciology, 58 (2012), pp. 889–903.
  - [51] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM Journal on Scientific Computing, 14 (1993), pp. 461–469.
  - [52] Y. SAAD, *Iterative methods for sparse linear systems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, second ed., 2003.

- [53] G. SCHUBERT, D. TURCOTTE, AND P. OLSEN, *Mantle Convection in the Earth and Planets*, Cambridge University Press, 2001.
- [54] H. SEROUSSI, H. B. E. N. DHIA, M. MORLIGHEM, E. LAROURE, E. RIGNOT, AND D. AUBRY, *Coupling ice flow models of varying orders of complexity with the Tiling method*, Journal of Glaciology, 58 (2012).
- [55] J. R. SHEWCHUK, *Triangle: Engineering a 2D quality mesh generator and delaunay triangulator*, in Applied Computational Geometry: Towards Geometric Engineering, M. C. Lin and D. Manocha, eds., vol. 1148 of Lecture Notes in Computer Science, Springer-Verlag, 1996, pp. 203–222. From the First ACM Workshop on Applied Computational Geometry.
- [56] I. K. TEZAUR, M. PEREGO, A. G. SALINGER, R. S. TUMINARO, AND S. F. PRICE, *Albany/felix: a parallel, scalable and robust, finite element, first-order Stokes approximation ice sheet solver built for advanced analysis*, Geoscientific Model Development, 8 (2015), pp. 1197–1220.
- [57] A. TOSELLI AND C. SCHWAB, *Mixed hp-finite element approximations on geometric edge and boundary layer meshes in three dimensions*, Numerische Mathematik, 94 (2003), pp. 771–801.
- [58] U. TROTTEMBERG, C. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press, London, 2001.
- [59] P. VANĚK, M. BREZINA, AND J. MANDEL, *Convergence of algebraic multigrid based on smoothed aggregation*, Numerische Mathematik, 88 (2001), pp. 559–579.
- [60] P. VANĚK, J. MANDEL, AND M. BREZINA, *Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems*, Computing, 56 (1996), pp. 179–196.
- [61] A. ZMITROWICZ, *Glaciers and laws of friction and sliding*, Acta mechanica, 166 (2003), pp. 185–206.